

NAKIVO® Backup & Replication v4 User Guide

Table of Contents

Table of Contents	2
Deployment.....	7
System Requirements.....	8
Deployment Requirements.....	8
Supported Hypervisors	9
Network Requirements.....	9
Feature Requirements	9
Solution Architecture	11
Transporter	11
Director	12
Backup Repository.....	12
Deployment Scenarios.....	13
Simple Deployment	13
Distributed Deployment	13
Using Direct SAN Access	14
Installing NAKIVO Backup & Replication	16
Installing on Windows	16
Installing on Linux.....	20
Troubleshooting Installation.....	24
Updating NAKIVO Backup & Replication	27
Updating on Windows	27
Updating on Linux.....	28
Uninstalling NAKIVO Backup & Replication	31
Uninstalling NAKIVO Backup & Replication on Windows	31
Uninstalling NAKIVO Backup & Replication on Linux	32
Editions	33
Enterprise Edition	33
Pro Edition.....	33
Enterprise for VMware Essentials Edition	33
Pro for VMware Essentials Edition	33

Free Edition	33
NAKIVO Backup & Replication Licensing	34
Licensing Model.....	34
Replacing a License.....	34
Configuration	35
Configuring General Settings.....	36
Changing Login and Password.....	36
Configuring Email Settings	36
Configuring Email Notifications.....	37
Configuring System Settings	38
Configuring Inventory.....	39
Adding vCenters and ESX(i) hosts.....	39
Managing Inventory	40
Configuring Transporters.....	43
About Transporters	43
Adding Installed Transporters.....	44
Viewing Transporter Details.....	45
Managing Transporters.....	46
Configuring Backup Repositories.....	49
About Backup Repositories.....	49
Adding Backup Repositories.....	50
Viewing Backup Repository Details	56
Managing Backup Repositories	59
Backup Repository How-Tos	68
Troubleshooting Backup Repositories	70
Creating and Managing Support Bundles	72
Creating Support Bundles	72
Managing Support Bundles.....	73
Viewing Events.....	74
Backup.....	75
Creating Backup Jobs.....	76
STEP 1: Choose VMs to Backup.....	77
STEP 2: Choose a Backup Repository.....	80
STEP 3: Specify Backup Schedule	83

STEP 4: Specify Backup Options	86
Viewing Backup Job Information.....	91
Backup Job Widgets.....	91
Managing VM Backups	95
Deleting Recovery Points	95
Deleting VM Backups.....	98
VM Backup How-Tos	101
How to Stage (Seed) Initial VM Backup	101
How to Back Up to Two Removable Disks in Rotation	101
How to Improve Backup Performance Over WAN	102
How to Stop Backing Up a VM	103
Troubleshooting VM Backup	104
Replication	107
Staging VM Replication	108
Creating Replication Jobs	110
STEP 1: Choose VMs to Replicate.....	111
STEP 2: Choose Replica Location.....	114
STEP 3: Specify Replication Schedule	117
STEP 4: Specify Replication Options	120
Viewing Replication Job Information.....	125
Replication Job Widgets.....	125
Troubleshooting VM Replication.....	129
Replica VM Cannot be Powered On	129
Recovery.....	130
Recovering Files from VM Backups.....	131
Opening the File Recovery Wizard	132
STEP 1: Choose a Backup and a Recovery Point.....	133
STEP 2: Choose Files to Recover.....	133
Troubleshooting File Recovery.....	136
Recovering Microsoft Exchange Objects.....	138
Opening the Exchange Object Recovery Wizard.....	138
STEP 1: Choose a Backup and a Recovery Point.....	139
STEP 2: Choose Application Items to Recover.....	139
Troubleshooting Exchange Objects Recovery.....	142

Recovering VMs from VM Backups.....	143
Opening the New Recovery Job Wizard	143
STEP 1: Choose VMs to Recover.....	144
STEP 2: Choose Recovery Location	144
STEP 3: Specify Recovery Options	145
Job Management.....	149
Job Alarms and Notifications	151
Running Jobs on Demand	151
Stopping Jobs	152
Editing Jobs	152
Disabling and Enabling Jobs	152
Grouping Jobs.....	153
Deleting Jobs	153
Cloud Integration	155
Amazon EC2 Integration	156
Setting up Integration with Amazon.....	156
Managing Amazon Integration.....	157
vCloud Director Integration.....	159
Backing Up and Replicating vCloud Director VMs.....	159
Recovering vCloud Director VMs back to vApps	159
Opening the Exchange Object Recovery Wizard.....	Error! Bookmark not defined.
STEP 1: Choose a Backup and a Recovery Point.....	Error! Bookmark not defined.
STEP 2: Choose Application Items to Recover.....	Error! Bookmark not defined.
Reporting.....	167
Setting up Email Reporting	168
Generating Reports on Demand	169
Command Line Interface	170
Using Command Line Interface.....	171
Using Command Line Interface Locally.....	171
Using Command Line Interface Remotely	171
Available Commands and Exit Codes.....	172
Available Commands	172
Exit Codes.....	173

Troubleshooting.....174

Generic Disk-Level Job Failure 175

Generic VM-Level Job Failure 176

VM Quiescing Failure.....177

Enabling CBT for a VM 178

Enabling SSH in ESXi Hosts..... 179

Attributions 180

Deployment

This section covers the following topics:

- [System Requirements](#)
- [Solution Architecture](#)
- [Deployment Scenarios](#)
- [Installing NAKIVO Backup & Replication](#)
- [Updating NAKIVO Backup & Replication](#)
- [Uninstalling NAKIVO Backup & Replication](#)
- [Editions](#)
- [NAKIVO Backup & Replication Licensing](#)

System Requirements

Deployment Requirements

Hardware Requirements

NAKIVO Backup & Replication can be installed on a virtual or physical machine with the following characteristics:

- Director and Onboard Transporter:
 - CPU: 2 CPUs
 - RAM: 4 GB RAM + 250 MB RAM for each concurrent job
 - Free space: 1GB
- Transporter:
 - CPU: 2 CPUs
 - RAM: 2 GB RAM + 250 MB RAM for each concurrent job
 - Free space: 1GB

Supported Operating Systems

NAKIVO Backup & Replication can be fully installed on the following operating systems:

- Windows:
 - Windows Server 2012 R2 Standard (x64)
 - Windows Server 2012 Standard (x64)
 - Windows Server 2008 R2 Standard (x64)
 - Windows 8 Professional (x86)
 - Windows 8 Professional (x64)
 - Windows 7 Professional (x86)
 - Windows 7 Professional (x64)
- Linux:
 - Ubuntu 12.04 Server (x64)
 - SUSE Linux Enterprise Server 11 SP2 (64-bit)
 - Red Hat Enterprise Linux 6.3 (64-bit)

Supported OS Localizations

NAKIVO Backup & Replication supports the following OS localizations:

- English
- Italian
- German
- French
- Spanish

Supported Web Browsers

NAKIVO Backup & Replication can be accessed from the following Web browsers:

- Google Chrome v26 or later

- Mozilla Firefox v21 or later

Supported Hypervisors

NAKIVO Backup & Replication supports the following hypervisors:

- VMware vSphere 5.5, 5.1, 5.0, and 4.1
- Standalone ESXi 5.5, 5.1, 5.0, and 4.1
- Standalone ESX 4.1

Full administrative permissions are required to perform backup, replication, and recovery tasks on supported hypervisors.

Network Requirements

NAKIVO Backup & Replication has been tested to work in the following minimal network conditions:

- **Latency (RTT):** up to 250 ms
- **Packet loss:** up to 1%
- **Bandwidth:** starting 1Mb/s

ICMP ping traffic should be allowed on all hosts where NAKIVO Backup & Replication components are installed as well as on all source and target ESX(i) hosts in order for the automatic Transporter selection feature to function properly.

Feature Requirements

VM Backup and Replication Requirements

NAKIVO Backup & Replication provides limited support for the following VM disk types:

- RDM disks are supported only in *virtual compatibility mode*. RDM disks in *physical compatibility mode* are not supported (since such disks are not affected by snapshots).
- Backed up RDM disks in virtual compatibility mode are recovered as thin disks.
- Replicas of source VMs with RDM disks in virtual compatibility mode are created with thin disks.

File Recovery Requirements

While NAKIVO Backup & Replication supports all VM file systems for backup, replication, and full VM recovery, only the following VM file systems are supported for *file recovery*:

- If a Transporter is installed on Windows:
 - NTFS
 - FAT32
- If a Transporter is installed on Linux:
 - NTFS
 - FAT32
 - EXT3
 - EXT4
 - XFS

- File recovery from Windows dynamic disks and Linux LVM volumes is available only in case the Transporter assigned to the backup repository is installed on a Linux OS.
- File recovery is not supported for independent, encrypted, and RAID disks.

Microsoft Exchange Objects Recovery Requirements

NAKIVO Backup & Replication supports the following versions of Microsoft Exchange for objects recovery:

- Microsoft Exchange 2010
- Microsoft Exchange 2007

Microsoft Exchange Logs Truncation Requirements

NAKIVO Backup & Replication supports the following versions of Microsoft Exchange for logs truncation:

- Microsoft Exchange 2013
- Microsoft Exchange 2010
- Microsoft Exchange 2007

Solution Architecture

NAKIVO Backup & Replication is a server application that is designed to run on a dedicated virtual or physical machine. All product components can be installed on a single machine or can be distributed across several machines. NAKIVO Backup and Replication consists of the two main components: Director and Transporter:




Refer to the sections below to learn more about the product components.

Transporter

A Transporter is an application that runs as a service on a Windows-based or Linux-based machine. The Transporter performs actual backup, replication, and recovery tasks, as well as handles compression, deduplication, and encryption.




A single Transporter can back up, replicate, and recover multiple VMs from multiple vCenters and ESX(i) hosts.



In most cases, only one Transporter per site is required. In large environments, multiple Transports can be deployed on a single site to distribute the data protection workload.

Director

A Director is an application that runs as a service on a Windows or Linux machine. The Director provides Web user interface and manages Transporters. A single Director can manage multiple Transporters, enabling the ability to achieve horizontal scalability.



Backup Repository

A backup repository is a folder used by NAKIVO Backup & Replication to store VM backups. By default, a backup repository is created with the Director installation. NAKIVO Backup & Replication supports the following repository types:

- CIFS (SMB) share
- Local folder on a machine where a Transporter is installed


Once added to the product, a backup repository can be selected in jobs.

Deployment Scenarios

NAKIVO Backup & Replication is a modular solution that can be fully installed on a single machine to protect small and mid-sized environments, as well as scale out horizontally and support large distributed environments with 10,000+ VMs. Refer to the sections below to learn more about the product deployment scenarios.

Simple Deployment


For a single site deployment, it is often sufficient to install both the Director and Transporter on a single VM within your infrastructure.



This deployment will enable the ability to back up, replicate, and recover multiple VMs from multiple ESX(i) hosts.

Distributed Deployment

If you have multiple sites and need to back up and/or replicate over WAN, install the Director and Transporter on one site, and at least one Transporter on all other sites.



Using Direct SAN Access


If you use a Fiber Channel or an iSCSI Storage Area Network (SAN), NAKIVO Backup & Replication can use direct SAN access for data retrieval. Using this storage access mode can significantly increase the speed of backup and replication, while decreasing the load on your production network.

To use the Direct SAN Access mode, make sure the following requirements are met:


- The source Transporter (the one which will read VM data from your SAN storage) must be installed on either of the following machines:
 - Physical machine which is connected to your SAN storage device
 - Virtual machine running on an ESX(i) host which is physically connected to your SAN storage device



- If you plan to install the source Transporter on a virtual machine, make sure that:
 - The ESX(i) host running the VM with the source Transporter is not managed by the vCenter which you plan to back up or replicate:



- The VM with the source Transporter is not located on the ESX(i) host which you plan to back up or replicate:



- If you are using iSCSI SAN, make sure the iSCSI initiator is configured on the machine where you plan to install the source Transporter.
- Make sure that your SAN's LUNs can be accessed from that machine where you plan to install the source Transporter.
- Make sure you log in to your SAN's LUNs from that machine where you plan to install the source Transporter.

Installing NAKIVO Backup & Replication

NAKIVO Backup & Replication can be fully installed on Windows and Linux operating systems. Note that in order to recover files from Linux-based VMs, the Transporter assigned to the backup repository from which you want to recover files should be installed on a Linux OS. Refer to the sections below to learn how to install the product on a Windows or Linux OS.


Installing on Windows

Installing the Full Solution on Windows

To install all product components (both Director and Transporter) on a Windows OS, run the NAKIVO Backup & Replication installer and follow the steps below:


STEP 1: Review license agreement

On the **License** step, read the license agreement. If you agree to the terms of the license, click **Accept**:



STEP 2: Set up installation options


Set up the installation options as described below:



- Leave the **Install full solution (Director and Transporter)** option selected in the Install type dropdown list. This way all product components will be installed.
- To change the default installation path, click **Change**, browse to the desired location, and click **OK**.
- To change the default port that is used to connect to the Web UI of NAKIVO Backup & Replication, enter a new value in the **Director Web HTTPS port** field. Make sure the port you specify is open in your firewall.
- To change the default port that is used by the Director to communicate with the Onboard Transporter, enter a new value in the **Transporter port** field. Make sure the port you specify is open in your firewall.
- To change the default ports that are used by the Onboard Transporter for data transfer, enter a new range in the **Data transfer ports** field. Make sure the ports from the range you specify are open in your firewall.
- To create a backup repository on the machine where NAKIVO Backup & Replication is installed, leave the **Create default backup repository** check box selected. To change the default backup repository location, click **Change**, browse to the desired location, and click **OK**.
- Click **Install**.

STEP 3: Complete the installation

1. After installation is complete, you can leave the **Open NAKIVO Backup & Replication in your browser** check box selected to start your default web browser with NAKIVO Backup & Replication web UI open when you click **Finish**.



2. Click **Finish**. You can log in to NAKIVO Backup & Replication by opening the following URL in your web browser:

`https://machine_IP_or_DNS:director_https_port`

By default, login and password are not required to access NAKIVO Backup & Replication. To prevent unauthorized access to the product, you can set up login credentials as described in the [Changing Login and Password](#) section.

Installing a Transporter on Windows

Transporter Installation Prerequisites

Prior to installing a Transporter, make sure the following prerequisites are met:

- Make sure the machine where you plan to install the Transporter has a connection to the relevant items below:
 - The machine where the Director is installed


- vCenter servers and ESX(i) hosts that you plan to back up or replicate (if you plan to retrieve VM data using the Transporter you are about to install)
 - Machines where you have installed other Transporters (if you plan to set up data transfer between an existing Transporter and the one you are about to install)
 - Backup repository (if you plan to assign the Transporter you are about to install to a backup repository)
 - vCenter and ESX(i) hosts to which you plan to replicate VMs (if you plan to write data to the target hosts and datastores using the Transporter you are about to install)
- If you have discovered vCenters and ESX(i) hosts using DNS names, make sure those DNS names can be resolved on the machine where you plan to install the Transporter.

Transporter Installation Steps

If you have already installed the full solution (both Director and Transporter) and want to deploy an additional Transporter, run the NAKIVO Backup & Replication installer and follow the steps below:


STEP 1: Review license agreement

On the **License** step, read the license agreement. If you agree to the terms of the license, click **Accept**.



STEP 2: Set up installation options

1. Choose **Install only Transporter** from the **Install type** dropdown list:




2. Do the following:
 - To change the default installation path, click **Change**, browse to the desired location, and click **OK**.

- To change the default port that is used to communicate with the Transporter, enter a new value in the **Transporter** port field. Make sure the port you specify is open in your firewall.

3. Click **Install**.

STEP 3: Complete the installation

1. After installation is complete, click **Finish**.



2. Add the Transporter to NAKIVO Backup & Replication as described in the [Adding Installed Transporters](#) section.

Installing on Linux

Linux Installation Prerequisites


In order to install and use NAKIVO Backup & Replication on a Linux OS, make sure the following requirements are met:

- On Ubuntu and SLES, NAKIVO Backup & Replication relies on the following packages:
 - cifs-utils
 - open-iscsi
 - ntfs-3g
- On RedHat Enterprise Linux, NAKIVO Backup & Replication relies on the following packages:
 - cifs-utils
 - iscsi-initiator-utils
 - ntfs-3g

Installing the Full Solution on Linux


Follow the steps below to install all components of NAKIVO Backup & Replication (both Director and Transporter) on a Linux OS:

1. Upload the installer file to the machine on which you want to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:
 - To upload the installer from a Windows-based machine, follow the steps below:
 - a. Download the free WinSCP client from <http://winscp.net>, install and run it.
 - b. Choose **SCP** from the **File protocol** dropdown menu.
 - c. Specify the IP address or the hostname of the Linux machine where you want to install the product in the **Host name** field.
 - d. Specify the username and password to the Linux machine in the appropriate fields.
 - e. Leave other fields as is and click **Login**.



- f. Click **Yes** in the message that appears.
- g. In the left view, find a folder that contains the Linux installer, in the right view, go up to the root folder.
- h. Drag and drop the installer from left to the right view.

- i. Choose **Binary** from the **Transfer settings** dropdown menu in the Copy dialog that appears.



- j. Click **Copy**.

- To upload the product from a Linux-based machine, run the following command:

```
wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'
```

2. Log in to the Linux machine and allow the execution of the installer file. For example:

```
chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh
```

3. Execute the installer file with root privileges. For example:

```
sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh
```

- Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press **Y** and then press **Enter**.
- Type **F** to install the full solution and press **Enter**.
- Specify the installation path for the product: Press **Enter** to accept the default installation path **"/opt/nakivo"** or enter a custom path and press **Enter**.
- Specify the Director HTTPS port (which will be used to access the Web UI of NAKIVO Backup & Replication): Press **Enter** to accept the default port **"4443"** or enter a custom port number and press **Enter**. Make sure the port you specify is open in your firewall.
- Specify the Transporter port (which will be used to connect to the Transporter that is installed by default with the Director): Press **Enter** to accept the default port **"9446"** or enter a custom port number (1 to 65535) and press **Enter**. Make sure the port you specify is open in your firewall.
- Specify a range of port numbers (from 1 to 65535) that will be used to transfer data by the Onboard Transporter. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
- Specify a path to the default backup repository: Press **Enter** to accept the default path **"/opt/nakivo/repository"** or enter a custom path and press **Enter** to begin installation process.

After the installation is complete, you can log in to NAKIVO Backup & Replication by opening the following URL in your web browser:

`https://machine_IP_or_DNS:director_https_port`

By default, login and password are not required to access NAKIVO Backup & Replication. To prevent unauthorized access to the product, you can set up credentials in Configuration.

Installing a Transporter on Linux

Transporter Installation Prerequisites


Prior to installing a Transporter, make sure the following prerequisites are met:

- Make sure the machine where you plan to install the Transporter has a connection to the relevant items below:
 - The machine where the Director is installed
 - vCenter servers and ESX(i) hosts that you plan to back up or replicate (if you plan to retrieve VM data using the Transporter you are about to install)
 - Machines where you have installed other Transporters (if you plan to set up data transfer between an existing Transporter and the one you are about to install)
 - Backup repository (if you plan to assign the Transporter you are about to install to a backup repository)
 - vCenter and ESX(i) hosts to which you plan to replicate VMs (if you plan to write data to the target hosts and datastores using the Transporter you are about to install)
- If you have discovered vCenters and ESX(i) hosts using DNS names, make sure those DNS names can be resolved on the machine where you plan to install the Transporter.

Transporter Installation Steps


If you have already installed the full solution (both Director and Transporter) and want to deploy an additional Transporter, run the NAKIVO Backup & Replication installer and follow the steps below:

1. Upload the installer file to the machine on which you want to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:
 - a. To upload the installer from a Windows-based machine, follow the steps below:
 - a. Download the free WinSCP client from <http://winscp.net>, install and run it.



- b. Choose **SCP** from the **File protocol** dropdown menu.

- c. Specify the IP address or the hostname of the Linux machine where you want to install the product in the **Host name** field.
- d. Specify the username and password to the Linux machine in the appropriate fields.
- e. Leave other fields as is and click **Login**.
- f. Click **Yes** in the message that appears.
- g. In the left view, find a folder that contains the Linux installer, in the right view, go up to the root folder.
- h. Drag and drop the installer from left to the right view.
- i. Choose **Binary** from the **Transfer settings** dropdown menu in the Copy dialog that appears.



- j. Click **Copy**.

- To upload the product from a Linux-based machine, run the following command:

```
wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'
```

- Allow the execution of the installer file. For example:

```
chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh
```

- Execute the installer file with root privileges. For example:

```
sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh
```

- Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press **Y** and then press **Enter**.
- Type **T** to install only the Transporter and press **Enter**.
- Specify the installation path for the product: Press **Enter** to accept the default installation path **"/opt/nakivo"** or enter a custom path and press **Enter**.
- Specify the Transporter port (used to connect to the Transporter): Press **Enter** to accept the default port **"9446"** or enter a custom port number and press **Enter** to begin installation process. Make sure the port you specify is open in your firewall.

After the installation is complete, add the Transporter to NAKIVO Backup & Replication as described in the [Adding Installed Transporters](#) section.

Troubleshooting Installation

Package Installation Failure on Ubuntu

If the “cifs-utils” or “open-iscsi”, or “ntfs-3g” packages have failed to install on Ubuntu OS, install the missing packages manually by following the steps below:

1. Update the Ubuntu packages by running the following command:

```
apt-get install update
```

2. After the update, do the following:

- a. To install the “cifs-utils” package, run the following command:

```
apt-get install cifs-utils
```

- b. To install the “open-iscsi” package, run the following command:

```
apt-get install open-iscsi
```

- c. To install the “ntfs-3g” package, run the following command:

```
apt-get install ntfs-3g
```

Package Installation Failure on RedHat Enterprise Linux

“cifs-utils” or “iscsi-initiator-utils” Installation Failure

If the “cifs-utils” or “iscsi-initiator-utils” packages have failed to install on RedHat Enterprise Linux OS, install the missing packages manually by following the steps below:

1. Connect the “rhel-server-6.3-x86_64-dvd.iso” image to the machine where NAKIVO Backup & Replication is installed.
2. Mount the ISO. For example:

- a. Go to the “media” folder:

```
cd /media
```

- b. Create the “cdrom” folder:

```
mkdir cdrom
```

- c. Mount the ISO:

```
mount /dev/cdrom /media/cdrom
```

3. Configure the “yum” installation manager to install the required packages from the ISO you have mounted:

- a. Edit “rhel-source.repo” file by running the following command:

```
vi /etc/yum.repos.d/rhel-source.repo
```

- b. Press the **Insert** key to open the file and add the following lines to the file:

```
[rhel-dvd]
name=Red Hat DVD
baseurl=file:///media/cdrom
```

```
enabled=1
gpgcheck=0
```

4. Save the file and exit the editor by pressing the **Escape** key, then typing:

```
:wq
```

5. Install the missing packages:

- a. To install the “cifs-utils” package, run the following command:

```
yum install cifs-utils
```

- b. To install the “iscsi-initiator-utils” package, run the following command:

```
yum install iscsi-initiator-utils
```

“ntfs-3g” Installation Failure

If the “ntfs-3g” package has failed to install on RedHat Enterprise Linux OS, install the package manually by following the steps below:

1. Download the “ntfs-3g” package to the machine where NAKIVO Backup & Replication is installed. For example, you can download the package from <http://pkgs.repoforge.org/fuse-ntfs-3g>
2. Install the downloaded package. For example:

```
rpm -ivh fuse-ntfs-3g-2013.1.13-2.el6.rf.x86_64.rpm
```

Package Installation Failure on SUSE Linux Enterprise Server

“cifs-utils” or “iscsi-initiator-utils” Installation Failure

If the “cifs-utils” or “open-iscsi” packages have failed to install on SUSE Linux Enterprise Server OS, install the missing packages by following the steps below:

1. Mount the 1st DVD disk named “SLES-11-SP2-DVD-x86_64-GM-DVD1.iso” to the VM where NAKIVO Backup & Replication is installed.
2. Mount the ISO. For example:

- a. Go to the “media” folder:

```
cd /media
```

- b. Create the “cdrom” folder:

```
mkdir cdrom
```

- c. Mount the ISO:

```
mount /dev/cdrom /media/cdrom
```

3. Do the following:

- To install the “cifs-utils” package, run the following command:

```
zypper in cifs-utils
```

- To install the “open-iscsi” package, run the following command:

```
zypper in open-iscsi
```

“ntfs-3g” Installation Failure

If the “ntfs-3g” package has failed to install on SUSE Linux Enterprise Server OS, install the package manually by following the steps below:

1. Download the “ntfs-3g” package to the machine where NAKIVO Backup & Replication is installed. For example, the package is available in the “SLED-11-SP2-DVD-x86_64-GM-DVD1.iso” disk in the “suse/x86_64” directory.
2. Install the downloaded package. For example:

```
rpm -ivh ntfs-3g-2011.4.12-0.7.8.x86_64.rpm
```

Product UI Does Not Open After Installation

If NAKIVO Backup & Replication UI does not open in your web browser, make sure the Director HTTPS port (4443 by default) is open in your firewall. You can turn off the firewall on the machine where the Director is installed as described below:

- On RedHat Enterprise Linux or SUSE Linux Enterprise Server, run the following command:

```
iptables -F
```

- On Ubuntu Linux, run the following command:

```
ufw disable
```

- On Windows, go to **Start > Control Panel > Windows Firewall**, click **Turn Windows firewall on or off**, and then select the **Turn off Windows firewall** option.

Updating NAKIVO Backup & Replication


Updating on Windows

To update your copy of NAKIVO Backup & Replication to a newer version, follow the steps below:


1. Download the latest product installer from www.nakivo.com.

NOTE: Trial installers do not update or replace licenses of the installed product.


2. Make sure no jobs are running in the product.
3. Run the installer on:
 - The machine where the Director is installed.
 - Each machine where you have additionally installed the Transporter.
4. On the **License** step, read the license agreement. If you agree to the terms of the license, click **Accept**.



5. Click **Update**.



6. When the update is complete, click **Finish**:



Updating on Linux

To update NAKIVO Backup & Replication on a Linux OS, you need to:


- Download the latest Linux installers from www.nakivo.com.
NOTE: Trial installers do not update or replace licenses of the installed product.
- Run the *full solution* installer on the machine where the Director is installed.
- If additional Transporters are deployed, run the installer on each machine where an additional Transporter is installed.

Refer to the sections below for installation instructions.


Updating Director and Transporter on Linux

To update the Director and/or Transporter to a newer version, follow the steps below:

1. Upload the Linux installer to the machine where the Director or Transporter is installed. Make sure you are using the *binary transfer mode*. For example:
 - To upload the installer from a Windows-based machine, follow the steps below:
 - a. Download the free WinSCP client from <http://winscp.net>, install, and run it.
 - b. Choose **SCP** from the **File protocol** dropdown menu.
 - c. Specify the IP address or the hostname of the Linux machine where you want to install the product in the **Host name** field.
 - d. Specify the username and password to the Linux machine in the appropriate fields.
 - e. Leave other fields as is and click **Login**.



- f. Click **Yes** in the message that appears.
- g. In the left view, find a folder that contains the Linux installer, in the right view, go up to the root folder.
- h. Drag and drop the installer from left to the right view.
- i. Choose **Binary** from the **Transfer settings** dropdown menu in the Copy dialog that appears.



- j. Click **Copy**.

- To upload the product from a Linux-based machine, run the following command:

```
wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'
```

2. Allow the execution of the installer file. For example:

```
chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh
```

3. Make sure no jobs are running in the product.
4. Execute the installer file with root privileges. For example:

```
sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh
```


5. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press “Y” and then press **Enter**.
6. Press the “U” key and then press **Enter** to confirm the update.
7. Press the “Y” and then press **Enter** to stop services and begin the update process.

Uninstalling NAKIVO Backup & Replication


Uninstalling NAKIVO Backup & Replication on Windows

To uninstall NAKIVO Backup & Replication, run the product uninstaller on the machine where the Director is installed, as well as on all machines where the Transporter is installed. Follow the step below to uninstall NAKIVO Backup & Replication:

1. Go to **Start > Control Panel** and run **Programs and Features**.
2. Select **NAKIVO Backup & Replication** and click **Uninstall**.
3. In the **NAKIVO Backup & Replication Uninstallation Wizard**, click **Uninstall**:



4. Click **Close** when the uninstallation process has completed:



Uninstalling NAKIVO Backup & Replication on Linux

Refer to the sections below to learn how to uninstall NAKIVO Backup & Replication on a Linux OS.

Uninstalling Director and Onboard Transporter on Linux

To uninstall the Director and Onboard Transporter (which is installed with the Director by default) on Linux, follow the steps below:

1. Run the “uninstall” script which is located in the “director” folder inside the product installation folder. If the product is installed in the default location, run:

```
/opt/nakivo/director/uninstall
```

2. Press “U” to confirm uninstallation and then press **Enter**.

Uninstalling Transporter on Linux

To uninstall the Transporter on Linux, follow the steps below:

1. Run the “uninstall” script which is located in the “transporter” folder inside the product installation folder. If the product is installed in the default location, run:

```
/opt/nakivo/transporter/uninstall
```

2. Press “U” to confirm uninstallation and then press **Enter**.

Editions

NAKIVO Backup & Replication is available in five editions. Refer to the sections below to learn more.

Enterprise Edition

NAKIVO Backup & Replication Enterprise Edition provides a complete data protection feature set including multi-tenancy, self-service, and dynamic license allocation. This edition supports all editions of VMware vSphere (from Essentials to Enterprise Plus).

Pro Edition

NAKIVO Backup & Replication Pro Edition provides a base data protection feature set and does not include multi-tenancy, self-service, and dynamic license allocation. This edition supports all editions of VMware vSphere (from Essentials to Enterprise Plus).

Enterprise for VMware Essentials Edition

Intended to help small businesses protect virtual infrastructure at an affordable price, NAKIVO Backup & Replication Enterprise for VMware Essentials Edition supports only Essentials and Essentials Plus editions of VMware vSphere. This edition has the same feature set as the Enterprise Edition and offers a complete data protection feature set including multi-tenancy, self-service, and dynamic license allocation.

Pro for VMware Essentials Edition

Intended to help small businesses protect virtual infrastructure at an affordable price, NAKIVO Backup & Replication Pro for VMware Essentials supports only Essentials and Essentials Plus editions of VMware vSphere. This edition has the same feature set as the Pro Edition.

Free Edition

NAKIVO Backup & Replication Free Edition provides a quick and simple way to make ad-hoc backups of running VMs for operational, archival, and portability purposes. To learn about differences between the Free and paid Editions, refer to the [NAKIVO Backup & Replication Comparison: Free vs. Paid](#) document. You can upgrade to the Pro or Essentials Edition by replacing the product license.

NAKIVO Backup & Replication Licensing

Licensing Model


NAKIVO Backup & Replication is licensed per physical CPU socket for each server on which you back up or replicate VMs. Licenses are required only for the source side of backup and replication, that is, you do not need to license target servers for your replicas or servers on which you want to recover backed up VMs.

By default, NAKIVO Backup & Replication includes a 15 day trial license for 6 sockets.

Replacing a License

To replace the current license, follow the steps below:

1. Click **Configuration** in the upper right corner of the application.
2. Go to the **Licensing** tab and click **Change License**.



3. Locate and open the license file in the window that appears.

Configuration

This section covers the following topics:


- [Configuring General Settings](#)
- [Configuring Inventory](#)
- [Configuring Transporters](#)
- [Configuring Backup Repositories](#)
- [Creating and Managing Support Bundles](#)
- [Viewing Events](#)

Configuring General Settings

Changing Login and Password

To change login and password used to access the product, follow the steps below:

1. Log in to NAKIVO Backup & Replication.
2. Click **Configuration** in the upper right corner of the product.
3. Go to the **General** tab and click **User account settings**.
4. In the title of the **User account settings** box, click **Edit**:




5. Fill out the **Login**, **Password**, and **Confirm Password** fields and click **Apply**.

Configuring Email Settings

Configure email settings so NAKIVO Backup & Replication can send email notifications as well as reports over email.

To configure email settings, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **General** tab and click **Email settings**.
3. In the title of the **Email settings** box, click **Edit**.



4. Fill out the fields and click **Send Test Email** to verify the settings are correct.
5. Click **Apply**.

After email settings are configured, the product will send the following emails to the recipient(s) specified in the **To** field:

- Notifications on errors
- Reports after each job completion


Email notification settings can be changed as described in the [Configuring Email Notifications](#) section.

Configuring Email Notifications

NAKIVO Backup & Replication can send notifications and reports over email.

To configure the automatic notifications, follow the steps below:

1. Click **Configuration** in the upper right corner of the product and go to the **General** tab.
2. Make sure you have configured email settings as described in the [Configuring Email Settings](#) section.
3. Click **Automatic reports**.
4. In the title of the **Automatic reports** box, click **Edit**:




5. Select or deselect the following email notification options:
 - **Send job reports on each job** – if this option is selected, NAKIVO Backup & Replication will send an HTML report after completion of every job (regardless of the job success or failure) to email addresses specified in the text field. Use a semi-colon to separate multiple email addresses.
***NOTE:** Changing this setting will affect only new jobs; existing jobs should be updated manually.*
 - **Send notifications on errors to** – if this option is selected, NAKIVO Backup & Replication will send an email notification on job, repository, infrastructure, connection, and other failures to email addresses specified in the text field. Use a semi-colon to separate multiple email addresses.
 - **Send notifications on warnings to** – if this option is selected, NAKIVO Backup & Replication will send an email notification on non-critical events, such as infrastructure change, to email addresses specified in the text field. Use a semi-colon to separate multiple email addresses.
 - **Send Overview report on schedule to job** – if this option is selected, NAKIVO Backup & Replication will generate the Overview report (which includes information about all jobs and groups in the product) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semi-colon to separate multiple email addresses.
6. Click **Apply**.

Configuring System Settings

To configure the system settings, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **General** tab and click **System settings**.
3. In the title of the **System settings** box, click **Edit**.



4. Select or deselect the following options:
 - **Store events for the last X days:** This option specifies the time period (from 10 to 365 days) during which the application events will be kept. Older events are automatically deleted.
 - **Automatically log out after X minutes of inactivity:** If this option is selected, the current user will be automatically logged out of the product after the specified period of inactivity.
 - **Retry failed jobs X times with Y minutes interval:** If this option is selected, failed VM backup, replication, and recovery will be automatically re-tried the specified number of times (from 2 to 10) and with the specified time interval (from 1 to 60). Jobs with failed VM backup, replication, and recovery remain in the “running” state until all retries either succeed or fail.
 - **Process every source VM only by one job at a time:** If this option is selected, each source VM will be processed by no more than one job at a time. Disable this option to allow processing a single VM by several jobs at a time (not recommended).


Configuring Inventory

Prior to creating backup, replication, or recovery jobs, you need to add your virtual infrastructure to the product's Inventory. The discovered infrastructure is added to the internal product database, which is updated every 1 hour by default.

Adding vCenters and ESX(i) hosts

To discover vCenters and standalone ESX(i) hosts, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Inventory** tab and click **Add New Server**:



3. In the **Hostname or IP** field, specify the hostname or IP address of the vCenter server or standalone ESX(i) host that you want to add to the inventory.

NOTE: vCenter-managed ESX(i) hosts should not be discovered directly by their IP addresses or hostnames. Instead, discover the vCenter that manages such ESX(i) hosts.

4. In the **Login** and **Password** fields, specify credentials of the vCenter server or standalone ESX(i) host that you want to add to the inventory.

NOTE: The credentials you specify should have full administrative privileges to the vCenter server or standalone ESX(i) host.

5. Specify the Web services port in the appropriate field and click **Discover**.


After discovery has completed successfully, you can exit Configuration and create a new job with the newly discovered VMs.

Managing Inventory

Editing Discovery Settings

If the hostname, IP address, credentials, or web services port of the discovered vCenter or ESX(i) host have changed, update the required fields in the product by following the steps below:

1. Click **Configuration** in the upper right corner of the product
2. Go to the **Inventory** tab.
3. Click the discovered vCenter or ESX(i) host you want to edit.
4. In the title of the discovered vCenter or ESX(i), click **Edit**:



5. Update the appropriate fields and click **Save and Update**.

Updating Infrastructure


NAKIVO Backup & Replication keeps information about the discovered infrastructure in its internal database, which is updated every 1 hour by default. During the inventory update, the product collects all required information about your virtual infrastructure, such as a list of hosts and VMs, their power state, and so on.

Only one vCenter server or standalone ESX(i) host can be updated at a time. Therefore, if you have discovered more than one vCenter server or standalone ESX(i) host, they will remain in queue until they are able to be updated.

Refer to the sections below to learn how to update the discovered infrastructure.

Changing Inventory Update Frequency

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Inventory** tab.
3. Do either of the following:
 - To prevent the product from automatically updating the inventory, deselect the **Update inventory every X hour(s)** check box.
 - To change the inventory update frequency, enter a new value in the **Update inventory every X hour(s)** field (from 0.1 to 48).




New settings are applied instantly and do not need to be saved.

Manually Updating All Inventory

To update all inventory items, follow the steps below:


1. Click **Configuration** in the upper right corner of the product and go to the **Inventory** tab.
2. Click **Update All**:



Manually Updating a Single vCenter or ESX(i) Host

To update a single vCenter or ESX(i) host, follow the steps below:


1. Click **Configuration** in the upper right corner of the product and go to the **Inventory** tab.
2. Click the discovered vCenter or ESX(i) host that you want to update.
3. In the title of the discovered vCenter or ESX(i), click **Update**:



Removing Discovered vCenters and ESX(i) Hosts from Inventory

To remove a discovered vCenter server or ESX(i) host from the inventory, follow the steps below:

1. Click **Configuration** in the upper right corner of the product and go to the **Inventory** tab.
2. Click the vCenter server or ESX(i) host that you want to remove from inventory.
3. In the title of the vCenter server or ESX(i), click **Remove**:



4. Click **Remove** in the message that appears.

NOTE: You will not be able to remove a vCenter server or ESX(i) host from the inventory if there is at least one job that backs up or replicates VMs on that vCenter server or ESX(i) host. To remove such server or host from the inventory, you first need to delete (or edit) the corresponding jobs so no VMs are backed up or replicated on the host or server being removed.


Configuring Transporters

About Transporters


A Transporter is an application that runs as a service on a Windows-based or Linux-based machine. When you install the full solution, a Transporter instance is installed with the Director. The default Transporter is called “Onboard Transporter”. The Transporter performs actual backup, replication, and recovery tasks, as well as handles compression, deduplication, and encryption.



A single Transporter can back up, replicate, and recover multiple VMs from multiple vCenters and ESX(i) hosts.



In most cases, only one Transporter per site is required. In large environments, multiple Transports can be deployed on a single site to distribute the data protection workload. If you plan to back up, replicate, or recover VMs over WAN, it is recommended to deploy at least one Transporter on each site.



If you plan to transfer data over WAN without a VPN connection from your source site to the target site, make sure the source and target Transporters are added to the product using external IP addresses or using DNS names that can be properly resolved in WAN (so the two Transporter can connect to each other).


Deploying multiple Transporters also enables the use of network acceleration and AES 256 encryption of traffic between a pair of Transporters.

Adding Installed Transporters

After you have installed a Transporter, you need to add the Transporter to the product so it can be used for backup, replication, and recovery tasks.

To add an installed Transporter to the product, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Transporters** tab and click **Add Transporter**:




3. Fill out the fields as described below:
 - Specify a name for the Transporter in the **Transporter name** field.
 - Specify the IP or hostname of the machine where the Transporter is installed in the **IP/Hostname** field.
NOTE: If you are adding the Transporter by a DNS name, make sure this DNS name can be resolved on the machines where the Director and any other Transporters (which you plan to use in conjunction with the current one) are installed.
 - Specify the port number that will be used to connect to the Transporter in the **Port** field.
 - Specify a range of port numbers (from 1 to 65535) that will be used to transfer data in the **Data transfer ports** field. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
 - Specify the maximum number of disks that the Transporter should process simultaneously in the **Maximum load** field.
 - Optionally, you can type a description in the **Description** field.
4. Click **Add**.

The Transporter is added to the product and can be used for backup, replication, and recovery.

Viewing Transporter Details

To view Transporter details, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Transporters** tab and click a Transporter.
3. The following information is displayed:




- **IP/Hostname:** The IP or hostname of the machine where the Transporter is installed.
- **Ports:** The port numbers that are used to connect to the Transporter.
- **Maximum load:** The maximum number of disks that the Transporter should process simultaneously.
- **Signature:** The signature of the Transporter's certificate.
- **Version:** The version number of the Transporter.
- **Description:** The Transporter description.

Managing Transporters

Editing Transporter Settings

To modify settings of an existing Transporter, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Transporters** tab and click a Transporter you want to edit.
3. In the Transporter title, click **Edit**:



4. Edit the following settings as required:
 - Specify a new name for the Transporter in the **Transporter name** field.
 - Specify a new IP or hostname of the machine where the Transporter is installed in the **IP/Hostname** field.
 - Specify a new port number that will be used to connect to the Transporter in the **Port** field.
 - Specify the new maximum number of disks that the Transporter should process simultaneously in the **Maximum load** field.
 - Optionally, you can type a new description in the **Description** field.
5. Click **Save**.

The changes you have made are applied to the Transporter.


Updating Transporters

By default, NAKIVO Backup & Replication updates the information about Transporters every hour. During the update, the product collects all required information about all Transporters.

Only one Transporter can be updated at a time therefore, if you have more than one Transporter, all others will remain in queue until they are able to be updated.

Changing Update Frequency of Transporters

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Transporters** tab.
3. Do either of the following:
 - To prevent the product from automatically updating all Transporters, deselect the **Update transporters every X [time period]** check box.
 - To change the update frequency, enter a new value in the **Update transporters every X [time period]** field (from 1 to 60 minutes or from 1 to 24 hours).




New settings are applied instantly and do not need to be saved.

Manually Updating All Transporters

To update all Transporters, follow the steps below:

1. Click **Configuration** in the upper right corner of the product and go to the **Transporters** tab.
2. Click **Update All**:




The update of all Transporters starts.

Manually Updating a Single Transporter

To update a single Transporter, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Transporters** tab.
3. Click the Transporter you want to update.
4. In the title of the Transporter, click **Update**:




The Transporter update starts.

Removing Transporters

To remove a Transporter from Nakivo Backup & Replication, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Transporters** tab.
3. Click the Transporter you want to remove.
4. In the Transporter title, click **Remove**:



NOTE: The following Transporters cannot be removed:

- The Onboard Transporter (which is installed with the Director by default).
- Transporters which are manually assigned to a job.

Configuring Backup Repositories


About Backup Repositories

A backup repository is a folder used by NAKIVO Backup & Replication to store VM backups. When you add a backup repository to the product, NAKIVO Backup & Replication creates a folder named “NakivoBackup” in the specified location and keeps all the VM data and backup repository metadata in that folder.


IMPORTANT: Do not modify or delete any files inside the “NakivoBackup” folder. Modifying or deleting any files inside the “NakivoBackup” folder may damage an entire backup repository.

The product is designed so that a single backup repository can store up to 16 TB of data. By default, VM backups are compressed and deduplicated at the block level across the entire backup repository to save storage space.

Each backup repository is managed by a single Transporter (called Assigned Transporter). In other words, only one Transporter can read data from and write data to a backup repository. The Assigned Transporter is responsible for all interaction with the backup repository. The Assigned Transporter of a repository can be changed.



Note that a single Transporter can be assigned to and manage multiple backup repositories:



You can add the following folders as a backup repository:

- Windows (CIFS) share
- Local folder on the machine where a Transporter is installed

By default, a backup repository is created when the full solution (both Director and Transporter) is installed. The default backup repository is named “Onboard repository”.

Adding Backup Repositories


Adding a Windows (CIFS) Share as a Backup Repository

Before creating a backup repository, make sure that:

- The folder where you want to create the backup repository already exists on the share.
- The share can be accessed from the machine where NAKIVO Backup & Replication is installed.
- Each backup repository is managed only by one Transporter (called Assigned Transporter) which writes data to and reads data from the backup repository. Make sure the share can be accessed from the machine where the Assigned Transporter is installed.
- You have credentials with the “write” permissions to the share.

To add a new backup repository or import an existing one on a Windows share, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab and click **Add Backup Repository**:



3. Fill out the fields as described below:
 - a. **Name:** Type in a name for the backup repository.
 - b. **Type:** Choose **Remote windows share** from the menu.
 - c. **Path to the share:** Specify a path to the share.
 - d. **Username:** Specify a username that will be used to access the share in the field.
 - e. **Password:** Specify a password that will be used to access the share in the field.
 - f. **Assigned Transporter:** From the dropdown menu, choose a Transporter that will be managing (i.e. writing data to and reading data from) this backup repository.
 - g. **Compress all backups in this repository:** Leave this option selected to keep all backups in a compressed form in order to save storage space. You can deselect this option if you are backing up to a hardware-based storage device (such as EMC Data Domain) that provides compression and deduplication.
***NOTE:** This option cannot be changed after the backup repository is created.*
 - h. **Deduplicate all backups in this repository:** Leave this option selected to store only unique blocks of data in the backup repository in order to save storage space. You can deselect this option if you are backing up to a storage device (such as EMC Data Domain) that provides compression and deduplication.
***NOTE:** This option cannot be changed after the backup repository is created.*
 - i. **Enable automatic repository self-healing:** Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure (such as incorrect timestamps on metadata and data files). You can deselect this option and run self-healing manually as described in the [Running Backup Repository Self-Healing](#) section.
 - j. **Run repository self-healing on schedule:** If required, select this check box to additionally run repository self-healing on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 am. For more information on self-healing, refer to the [Running Backup Repository Self-Healing](#) section.

- k. **Run full data verification on schedule:** If selected, the product will run full verification of all data available in the backup repository on the specified schedule. The product will read each block of data and ensure that it is identical to the data block that was read on the source VM during the backup. This way the product will verify each recovery points in the backup repository.


NOTE: Backup verification is a time-consuming process and consumes CPU of the Transporter assigned to the backup repository. It is recommended to schedule backup verification during non-working times.

- l. **Reclaim unused space on schedule:** If required, select this option to run the backup repository space reclaim process on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every Saturday at 12 pm. For more information on reclaiming free space, refer to the [Verifying VM Backups](#)
- m. **The backup verification process reads each block of data inside the backup repository and ensures that it is identical to the data block that was read on the source VM during the backup.** This way the product verifies that files, application objects, and entire VMs can be recovered from all VM backups.

NOTE: Before the backup verification begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the backup verification process is in progress.


To verify all VM backups in a repository, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the Backup **repositories** tab and click a backup repository.
3. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The backup verification process cannot be started if a job which backs up to this backup repository is running.

In the dialog box that appears, select the **Perform full data verification** option and click **Start**:



The backup verification process is started.

Dealing with Corrupted Recovery Points

If the backup verification process finds that a recovery point is corrupted, do the following:


- If the last recovery point is corrupted, follow the steps below to repair it:
 - a. Delete the backup job (as described in the [Deleting Jobs](#) section) or edit the backup job (as described in the [Editing Jobs](#) section) and remove the source VM that has the corrupted recovery point.
 - b. Create a new backup job and map the source VM to the backup with the corrupted recovery point as described in the [Mapping Source VMs to Existing Backups](#) section.
 - c. Run the job. The recovery point will be repaired using the data from the source VM.
- If the corrupted recovery point is not the last one, it cannot be repaired. Delete such recovery point.

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the backup verification process, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the Backup **repositories** tab and click a backup repository.
3. In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The backup verification process is stopped.


- d. Reclaiming Backup Repository Space section.
 - e. **Description:** Optionally, type a backup repository description in the field.
4. Click **Add**.

The backup repository is created with the options that you have specified. If a backup repository already exists in the specified location, the backup repository is imported in NAKIVO Backup & Replication.

Adding a Local Folder as a Backup Repository

To add a new backup repository or to import an existing one, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab and click **Add Backup Repository**:



3. Fill out the fields as described below:

- a. **Name:** Type in a name for the backup repository.
- b. **Type:** Choose **Local folder on assigned transporter** from the dropdown list.
- c. **Assigned transporter:** Choose a Transporter that is installed on the machine where you want to create a backup repository. This Transporter will be managing (i.e. writing data to and reading data from) this backup repository.
- d. **Path to the local folder:** Specify a local path (on the machine where the Assigned Transporter is installed) where you want to create a backup repository. For example, if the Transporter is installed on a Linux OS, you can specify the following path:

`/opt/nakivo/repository`

- e. **Compress all backups in this repository:** Leave this option selected to keep all backups in a compressed form in order to save storage space. You can deselect this option if you are backing up to a hardware-based storage device (such as EMC Data Domain) that provides compression and deduplication.

NOTE: This option cannot be changed after the backup repository is created.

- f. **Deduplicate all backups in this repository:** Leave this option selected to store only unique blocks of data in the backup repository in order to save storage space. You can deselect this option if you are backing up to a storage device (such as EMC Data Domain) that provides compression and deduplication.

NOTE: This option cannot be changed after the backup repository is created.

- g. **Enable automatic repository self-healing:** Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure (such as incorrect timestamps on metadata and data files). You can deselect this option and run self-healing manually as described in the [Running Backup Repository Self-Healing](#) section.
- h. **Run repository self-healing on schedule:** If required, select this check box to additionally run repository self-healing on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 am. For more information on self-healing, refer to the [Running Backup Repository Self-Healing](#) section.
- i. **Run full data verification on schedule:** If selected, the product will run full verification of all data available in the backup repository on the specified schedule. The product will read each block of data and ensure that it is identical to the data block that was read on the source VM during the backup. This way the product will verify each recovery points in the backup repository.


NOTE: Backup verification is a time-consuming process and consumes CPU of the Transporter assigned to the backup repository. It is recommended to schedule backup verification during non-working times.

- j. **Reclaim unused space on schedule:** If required, select this option to run the backup repository space reclaim process on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every Saturday at 12 pm. For more information on reclaiming free space, refer to the [Verifying VM Backups](#)
- k. **The backup verification process reads each block of data inside the backup repository and ensures that it is identical to the data block that was read on the source VM during the backup. This way the product verifies that files, application objects, and entire VMs can be recovered from all VM backups.**

NOTE: Before the backup verification begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the backup verification process is in progress.


To verify all VM backups in a repository, follow the steps below:

4. Click **Configuration** in the upper right corner of the product.
5. Go to the Backup **repositories** tab and click a backup repository.
6. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The backup verification process cannot be started if a job which backs up to this backup repository is running.

In the dialog box that appears, select the **Perform full data verification** option and click **Start**:



The backup verification process is started.

Dealing with Corrupted Recovery Points

If the backup verification process finds that a recovery point is corrupted, do the following:


- If the last recovery point is corrupted, follow the steps below to repair it:
 - l. Delete the backup job (as described in the [Deleting Jobs](#) section) or edit the backup job (as described in the [Editing Jobs](#) section) and remove the source VM that has the corrupted recovery point.
 - m. Create a new backup job and map the source VM to the backup with the corrupted recovery point as described in the [Mapping Source VMs to Existing Backups](#) section.
 - n. Run the job. The recovery point will be repaired using the data from the source VM.
- If the corrupted recovery point is not the last one, it cannot be repaired. Delete such recovery point.

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the backup verification process, follow the steps below:

7. Click **Configuration** in the upper right corner of the product.
8. Go to the Backup **repositories** tab and click a backup repository.
9. In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The backup verification process is stopped.

- o. Reclaiming Backup Repository Space section.
- p. **Description:** Optionally, type a backup repository description in the field.


10. Click **Add**.

The backup repository is created with the options that you have specified. If a backup repository already exists in the specified location, the backup repository is imported in NAKIVO Backup & Replication.

Viewing Backup Repository Details

To view backup repository details, follow the steps below:

1. Click **Configuration** in the upper right corner of the product and go to the **Backup repositories** tab.
2. Click a backup repository.
3. The following data is displayed:




- **Free:** The amount of free space currently available on the backup repository.
- **Used:** The amount of space that the backup repository occupies on a disk. The amount of space that can be reclaimed is displayed in parentheses. You can learn more about the space reclaim process in the [Verifying VM Backups](#)
- The backup verification process reads each block of data inside the backup repository and ensures that it is identical to the data block that was read on the source VM during the backup. This way the product verifies that files, application objects, and entire VMs can be recovered from all VM backups.

NOTE: Before the backup verification begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the backup verification process is in progress.


To verify all VM backups in a repository, follow the steps below:

4. Click **Configuration** in the upper right corner of the product.
5. Go to the **Backup repositories tab** and click a backup repository.
6. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The backup verification process cannot be started if a job which backs up to this backup repository is running.

In the dialog box that appears, select the **Perform full data verification** option and click **Start**:



The backup verification process is started.

Dealing with Corrupted Recovery Points

If the backup verification process finds that a recovery point is corrupted, do the following:


- If the last recovery point is corrupted, follow the steps below to repair it:
 - q. Delete the backup job (as described in the [Deleting Jobs](#) section) or edit the backup job (as described in the [Editing Jobs](#) section) and remove the source VM that has the corrupted recovery point.
 - r. Create a new backup job and map the source VM to the backup with the corrupted recovery point as described in the [Mapping Source VMs to Existing Backups](#) section.
 - s. Run the job. The recovery point will be repaired using the data from the source VM.
- If the corrupted recovery point is not the last one, it cannot be repaired. Delete such recovery point.

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the backup verification process, follow the steps below:

7. Click **Configuration** in the upper right corner of the product.
8. Go to the Backup **repositories** tab and click a backup repository.
9. In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The backup verification process is stopped.

- Reclaiming Backup Repository Space section.
- **Compression:** The backup compression ratio (i.e. the percentage of actual data written to the backup repository after compression. For example, if 200 GB of data is compressed down to 50 GB, the compression ratio is calculated as 25%).
- **Deduplication:** The data deduplication ratio. The first number represents the number of deduplicated blocks (those which are stored once but are used in multiple backups). The second number represents the number of unique blocks.
- **Automatic self-healing:** The current state of the automatic self-healing option for the backup repository.
- **Scheduled self-healing:** The current state of the scheduled self-healing option for the backup repository.
- **Scheduled space reclaiming:** The current state of the scheduled space reclaiming option for the backup repository.
- **Type:** The type of backup repository, which can be either of the following:
 - Remote Windows Share
 - Local folder on assigned Transporter
- If the backup repository type is **Local folder on assigned Transporter**, the local path to the backup repository folder on a disk is displayed.
- **Assigned Transporter:** The Transporter that manages the backup repository (i.e. reads data from and writes data to the backup repository).
- **Description:** The optional backup repository description.
- **Backups:** List of available backups in the backup repository

Managing Backup Repositories


Verifying VM Backups

The backup verification process reads each block of data inside the backup repository and ensures that it is identical to the data block that was read on the source VM during the backup. This way the product verifies that files, application objects, and entire VMs can be recovered from all VM backups.

NOTE: Before the backup verification begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the backup verification process is in progress.


To verify all VM backups in a repository, follow the steps below:

10. Click **Configuration** in the upper right corner of the product.
11. Go to the **Backup repositories** tab and click a backup repository.
12. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The backup verification process cannot be started if a job which backs up to this backup repository is running.

In the dialog box that appears, select the **Perform full data verification** option and click **Start**:



The backup verification process is started.

Dealing with Corrupted Recovery Points

If the backup verification process finds that a recovery point is corrupted, do the following:

- If the last recovery point is corrupted, follow the steps below to repair it:
 - t. Delete the backup job (as described in the [Deleting Jobs](#) section) or edit the backup job (as described in the [Editing Jobs](#) section) and remove the source VM that has the corrupted recovery point.


- u. Create a new backup job and map the source VM to the backup with the corrupted recovery point as described in the [Mapping Source VMs to Existing Backups](#) section.
- v. Run the job. The recovery point will be repaired using the data from the source VM.
- If the corrupted recovery point is not the last one, it cannot be repaired. Delete such recovery point.

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the backup verification process, follow the steps below:


13. Click **Configuration** in the upper right corner of the product.
14. Go to the **Backup repositories** tab and click a backup repository.
15. In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The backup verification process is stopped.

Reclaiming Backup Repository Space

When a backup or a recovery point is deleted in a backup repository, the space occupied by that backup or recovery point is marked as “free” and can be reused by new data blocks on the next job runs. However, the actual size of the backup repository may not change. The size of a backup repository can be reduced by rearranging the data blocks so there are no “free” ones that occupy storage space. The amount of space that can be freed up is displayed in parentheses after the amount of used space:




Reclaiming free space can take the same amount of time as copying the entire backup repository to the storage where it is located (that is, if your repository size is 500 GB, reclaiming free space can take the same amount of time as copying 500GB of data to the storage where the backup repository is located).

NOTE: Before the space reclaim process begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the space reclaim process is in progress. If you need to run a backup job before backup the space reclaim process has completed, you can do so as described in the [Stopping the Space Reclaim Process](#) section.


To reclaim free space, follow the steps below:

16. Click **Configuration** in the upper right corner of the product.
17. Go to the **Backup repositories** tab and click a backup repository.
18. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The space reclaim process cannot be started if a job which backs up to this backup repository is running.

19. In the dialog box that appears, leave the **Reclaim unused storage space** check box selected and click **Start**:



The process of rearranging data blocks is started, and progress is displayed in the title of the backup repository.

Stopping the Space Reclaim Process


You can stop the space reclaim process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

NOTE: Before the space reclaim begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the space reclaim process is in progress.

To stop the space reclaim process, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.

- Go to the **Backup repositories** tab and click a backup repository.
- In the title of the backup repository, click **Manage** and then click **Stop maintenance**: bgfg



The space reclaim process is stopped.

Running Backup Repository Self-Healing


The self-healing process verifies backup repository integrity and automatically repairs errors if possible. Namely, the process performs the following tasks:

- Verifies that data blocks of each recovery point are present in the backup repository
- Cleans up “in progress” blocks of data from failed/crashed backup job runs that had no proper cleanup
- Verifies and repairs backup repository metadata so that it correctly describes available data
- Restores the consistent state of the backup repository to enable subsequent backup jobs


NOTE: Before the self-healing process begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the self-healing process is in progress. Stopping the process is described in the [Stopping the Self-Healing Process](#) section.

To run the backup repository self-healing, follow the steps below:

- Click **Configuration** in the upper right corner of the product.
- Go to the **Backup repositories** tab and click a backup repository.
- In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



- In the dialog box that appears, the **Perform repository self-healing** option is already selected. Click **Start**:




The self-healing process is started.

Stopping the Self-Healing Process

You can stop the self-healing process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the self-healing process, follow the steps below:

- Click **Configuration** in the upper right corner of the product.
- Go to the **Backup repositories** tab and click a backup repository.
- In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The self-healing process is stopped.


Detaching Backup Repositories

Detaching a backup repository saves the backup repository data and metadata in a consistent state and then stops the product's interaction with the backup repository (such as read and write of data and metadata, and so on). You may want to detach a backup repository to move it to a different location or to put the storage with the backup repository on maintenance.

IMPORTANT: Since the product stops working with detached backup repositories, jobs that back up VMs to a detached backup repository will fail.

To detach a backup repository, follow the steps below:

- Click **Configuration** in the upper right corner of the product.
- Go to the **Backup repositories** tab and click a backup repository.
- In the backup repository title, click **Manage** and then click **Detach**:




NOTE: A backup repository cannot be detached if a job which backs up to this backup repository is running.

The backup repository is detached from the product. You can reattach the backup repository to NAKIVO Backup & Replication as described in the [Attaching Backup Repositories](#) section.

Attaching Backup Repositories

If you have detached a backup repository you can reattach it to the product by following the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab and click a backup repository.
3. In the backup repository title, click **Manage** and then click **Attach**:




The backup repository is reattached to NAKIVO Backup & Replication. You can now back up to the attached backup repository.

Editing Backup Repositories

To modify settings of an existing backup repository, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab and click a backup repository.
3. In the backup repository title, click **Manage** and then click **Edit**:



NOTE: A backup repository cannot be edited if there is a job which backs up to this backup repository is running.

4. Update the fields as necessary:

NOTE: The **Compress backups in this repository** and **Deduplicate backups in this repository** options cannot be changed.

- **Name:** Update the name of the backup repository if necessary.
- **Type:**
 - If the **Remote windows share** option is selected:
 - Specify a new path to the share in the **Path to the share** field.
 - Specify a new username that will be used to access the share in the **Username** field.
 - Specify a new password that will be used to access the share in the **Password** field.
 - From the **Assigned Transporter** dropdown menu, choose a new Transporter that will be managing (i.e. writing data to and reading data from) this backup repository.
 - If the **Local folder on assigned transporter** option is selected:
 - From the **Assigned Transporter** dropdown menu, choose a new Transporter that will be managing (i.e. writing data to and reading data from) this backup repository.
 - Specify a new local path (on the machine where the Assigned Transporter is installed) to the backup repository in the **Path to the local folder** field.
- **Enable automatic repository self-healing:** Update the setting if necessary.
- **Run repository self-healing on schedule:** Update the settings if necessary.
- **Reclaim unused space on schedule:** Update the settings if necessary.
- **Description:** Update repository description if necessary.

5. Click **Save**.

Changes you have made are applied and the backup repository update starts.


Updating Backup Repositories

By default, NAKIVO Backup & Replication updates the information about backup every 1 hour. During the update, the product collects all required information about backup repositories (such as the amount of free space, number of backups and recovery points, and so on).

Only one backup repository is updated at a time. Therefore, if you have more than one backup repository, all others will remain in queue until they are able to be updated.

Changing Update Frequency of Backup Repositories

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab.
3. Do either of the following:
 - To prevent the product from automatically updating all backup repositories, deselect the **Update backup repositories every X [time period]** check box.
 - To change the update frequency, enter a new value in the **Update backup repositories every X [time period]** fields (from 1 to 60 minutes or from 1 to 24 hours).




New settings are applied instantly and do not need to be saved.

Manually Updating All Backup Repositories

To update all backup repositories, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab.
3. Click **Update All**:




Backup repositories update starts.

Manually Updating a Single Backup Repository

To update a single backup repository, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab.
3. Click the backup repository that you want to update.
4. In the title of the backup repository, click **Update**:



Backup repository update starts.


Removing Backup Repositories

In NAKIVO Backup & Replication, you can either permanently delete a backup repository and all of its data or remove only a backup repository from the product while keeping all of its data. After removing a backup repository you will be able to import it into the same or a new instance of the product.

NOTE: You will not be able to remove a backup repository if there is a job that backs up VMs to this backup repository. To remove a backup repository, you first need to delete (or edit) the corresponding jobs so no VMs are backed up to the backup repository being removed.

To permanently delete or remove a backup repository from the product, follow the steps below:

1. Click **Configuration** in the upper right corner of the product
2. Go to the **Backup repositories** tab.
3. Click a backup repository.
4. In the backup repository title, click **Manage** and then click **Remove**:



5. Do the following when the confirmation message appears:
 - To remove the backup repository from NAKIVO Backup & Replication and keep the backup repository on a disk, click the **Remove Repository and Keep Backups** button.
- NOTE:** You can add the removed backup repository back to the same or to a new product installation as described in the [Adding Backup Repositories](#) section.
- To permanently delete the backup repository and all its data, click the **Remove Repository and Delete Backups** button.


IMPORTANT: This operation will permanently delete the backup repository and all VM backups.

Backup Repository How-Tos

How to Back Up to an NFS, NAS, iSCSI, or Other Storage Type

To back up to an NFS, NAS, or iSCSI or other storage type, follow the steps below:

1. Mount your storage to the machine where a Transporter is installed.
2. Create a new backup repository on the mounted storage as described below:
 - a. Click **Configuration** in the upper right corner of the product.
 - b. Go to the **Backup repositories** tab and click **Add Backup Repository**:



- c. Fill out the fields as follows:
 - Type in a name for the backup repository in the **Name** field.
 - Choose **Local folder on assigned transporter** from the **Type** menu.
 - From the **Assigned Transporter** dropdown menu, choose the Transporter to which you have mounted your storage.
 - Specify a local path (on the machine where the Transporter is installed) where you want to create a backup repository in the **Path to the local folder** field.
 - Leave the **Compress all backups in this repository** option selected to keep all backups in a compressed form in order to save storage space. You can deselect this option if you are backing up to a hardware-based storage device (such as EMC Data Domain) that provides compression and deduplication.
NOTE: This option cannot be changed after the backup repository is created.
 - Leave the **Deduplicate all backups in this repository** option selected to store only unique blocks of data in the backup repository in order to save storage space. You can deselect this option if you are backing up to a storage device (such as EMC Data Domain) that provides compression and deduplication.
NOTE: This option cannot be changed after the backup repository is created.
 - Optionally, you can type a backup repository description in the **Description** field.

3. Click **Add**.

The backup repository is created with the options you have specified. You can now choose the newly created backup repository when creating new backup jobs.

How to Move a Backup Repository to a New Location

To move a backup repository to a new location, follow the steps below:

1. Make sure no jobs which back up to the backup repository you want to move are running.
2. To ensure data consistency, detach the backup repository as described in the [Detaching Backup Repositories](#) section.
3. Move the "NakivoBackup" folder (which is created in the location you specify during backup repository creation) and all of its contents to a new location.

IMPORTANT: Do not modify or delete any files inside the “NakivoBackup” folder. Modifying or deleting any files inside the “NakivoBackup” folder may damage an entire backup repository.

4. If the path or access credentials to the backup repository have changed, edit the backup repository as described in the [Editing Backup Repositories](#) section and specify correct information.
5. Attach the backup repository to the product as described in the [Attaching Backup Repositories](#) section.

How to Import a Backup Repository

If you have lost the machine with NAKIVO Backup & Replication but still have a backup repository, you can install a new instance of the product and then import your existing backup repository in that instance. Simply follow the steps specified in the [Attaching Backup Repositories](#) section and specify a correct location to your existing backup repository.

NOTE: Backup job will not be recreated when importing an existing backup repository. However, you can recover files and VMs from the imported backup repository, as well as back up VMs to the imported backup repository.

How to Back Up to a Linux Server

To back up to a Linux server, follow the steps below:

1. Install a Transporter on the Linux machine where you want to keep your backups as described in the [Installing a Transporter on Linux](#) section.
2. Add the newly installed Transporter to the product as described in the [Adding Installed Transporters](#) section.
3. Create a new backup repository on the machine where you have installed the Transporter as described in the [Adding a Local Folder as a Backup Repository](#) section.

Now you can choose the newly created backup repository when creating new backup jobs.

Troubleshooting Backup Repositories

A Problem with a Backup Infrastructure

Network connectivity disruptions, change of credentials and permissions, improper backup repository detach, and even third-party software can cause metadata and backup data to be locked, partially written, or even corrupted in a backup repository. To ensure data integrity at all times, the product periodically runs quick self-checks on backup repositories.

In case a quick self-check finds symptoms of a problem in a backup infrastructure – such as incorrect timestamps on metadata and data files – the backup repository is locked to prevent any potential data corruption. To unlock the backup repository, run the full self-check as described in the [Running Backup Repository Self-Healing](#) section. If this problem persists, please refer to NAKIVO Technical Support for further investigation.

Backup Repository on a Linux VM Cannot be Added due to Insufficient Permissions

In case you are trying to add a backup repository on a Linux OS and get a message that a Transporter cannot read and/or write to the specified location, do the following:

1. Make sure that the folder where you want to add the backup repository is managed by the “bhsvc” user. To assign the “bhsvc” user to the folder, you can run the following command:

```
chown -R bhsvc:bhsvc /path_to_the_folder
```

2. In NAKIVO Backup & Replication, add the backup repository.
3. After the backup repository is added, make sure that repository files and folders have correct permissions.

Here’s how:

- a. Navigate to the backup repository folder. For example:

```
cd /path_to_the_folder/NakivoBackup
```

- b. List directory contents, by running the following command:

```
ll
```

- c. Make sure that the “transit” and “raw” folders have the following permissions:

```
owner rwx
```

```
group rwx
```

```
others ---
```

In case permissions are different, you can set the right permissions by running the following command:

```
chmod 77 /path_to_the_folder/NakivoBackup/folder_name
```

- d. Make sure that the “.desc” and “.desc2” files have the following permissions:

```
owner rw-
```

```
group r--
```

```
others ---
```

- e. In case permissions are different, you can set the right permissions by running the following command:

```
chmod 64 /path_to_the_folder/NakivoBackup/file_name
```

Backup Repository Cannot be Created on a Share

Make sure that:

- The folder where you want to create the repository already exists on the share.
- The share can be accessed from the machine where NAKIVO Backup & Replication is installed.
- Each backup repository is managed only by one Transporter (called Assigned Transporter) which writes data to and reads data from the backup repository. Make sure the share can be accessed from the machine where the Assigned Transporter is installed.
- You have credentials with the “write” permissions to the share.

Creating and Managing Support Bundles


To speed up investigation of technical support issues, Nakivo Backup & Replication provides the ability to create support bundles – a zipped collection of the product logs and system information. Sending a support bundle to the NAKIVO Support Team allows them to quickly identify the root cause of issues and to suggest a proper solution.

Refer to the sections below to learn how to create and manage support bundles.

Creating Support Bundles

To create a support bundle, follow the steps below:

1. Click **Support** in the upper right corner of the product.
2. Click **Create Support Bundle** in the window that appears.




3. Explain the problem in the **Please describe the problem you're experiencing** field.
4. Do either of the following:
 - To create a support bundle and immediately send it to NAKIVO Support Team, enter your email address in the **Contact email** field and click the **Create & Email to Support** button. You should receive an answer from NAKIVO Support Team within one business day.
NOTE: *Email settings should be configured as described in the [Configuring Email Settings](#) section.*
 - To create a support bundle that you can later download, click the **Create Support Bundle** button.

Managing Support Bundles

Downloading Support Bundles

To download a support bundle that you have previously created, follow the steps below:

1. Click **Support** in the upper right corner of the product.
2. Hover the mouse pointer over the title of a support bundle.
3. Click **Download**:




The support bundle will be sent to your web browser for download.

Removing Support Bundles

To remove a support bundle from the list, follow the steps below:

1. Click **Support** in the upper right corner of the product.
2. Hover the mouse pointer over the title of a support bundle.
3. Click **Remove**:



The support bundle is removed from the product.

Viewing Events

NAKIVO Backup & Replication can store and display system events. By default, events are stored for 60 days; you can specify the time period during which events are stored as described in the [Configuring System Settings](#) section.

To view events, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Events** tab.
3. If required, navigate between pages by using the **Page** control:

The screenshot shows the NAKIVO Backup & Replication Configuration page with the Events tab selected. The left sidebar contains navigation links: General, Inventory, Transporters, Backup repositories, Cloud integration, Events, and Licensing. The main content area displays a list of events. The 'Page' control at the bottom left is circled in red, showing '1' of 1951. The 'Exit Configuration' button is at the bottom right.

Event	Time
Amazon EC2 cloud integration was updated Updated Amazon EC2 cloud integration "NAKIVOCriticalVMs", time spent 8 seconds.	Jan 15, 2014 at 11:11 AM
Transporter was updated Transporter "NAKIVOCriticalVMs" was updated, time spent: 0 seconds.	Jan 15, 2014 at 11:11 AM
Transporter update has started Update has started on "NAKIVOCriticalVMs" transporter. 100%	Jan 15, 2014 at 11:11 AM
Transporter was updated Transporter "NAKIVOCriticalVMs" was updated, time spent: 0 seconds.	Jan 15, 2014 at 11:11 AM
Transporter update has started Update has started on "NAKIVOCriticalVMs" transporter. 100%	Jan 15, 2014 at 11:11 AM
Amazon EC2 cloud pricing retrieval is complete The latest prices are: Linux on-demand micro instance per hour in 0.001 USD: {EU1=20, SAMERO1=27, ASIA3=20, US_EAST1=20, US_WEST1=25, ASIA1=20, ASIA2=27, US_WEST2=20} Standard EBS per GB per month in 0.001 USD: {EU1=110, SAMERO1=190, ASIA3=110, US_EAST1=100, US_WEST1=110, ASIA1=110, ASIA2=120, US_WEST2=100}	Jan 15, 2014 at 11:11 AM
Transporter was updated Transporter "Onboard transporter" was updated, time spent: 0 seconds.	Jan 15, 2014 at 11:11 AM
Transporter was updated Transporter "NAKIVOCriticalVMs" was updated, time spent: 0 seconds.	Jan 15, 2014 at 11:11 AM

Page 1 of 1951

Displaying 1-20 of 39005

Exit Configuration

© 2014 NAKIVO, Inc. All rights reserved


Backup

This section covers the following topics:

- [Creating Backup Jobs](#)
- [Viewing Backup Job Information](#)
- [Managing VM Backups](#)
- [VM Backup How-Tos](#)
- [Troubleshooting VM Backup](#)

Creating Backup Jobs

In NAKIVO Backup & Replication, you can back up VMs by creating a job that specifies which VMs should be backed up, where the backups should be located, how often the backup should occur, and what backup options should be used. To create a backup job, click **Create** and then click **Backup Job** on any dashboard:



The New Backup Job Wizard opens. Complete the wizard as described in the sections below:


- STEP 1: Choose VMs to Backup
- STEP 2: Choose a Backup Repository
- STEP 3: Specify Backup Schedule
- STEP 4: Specify Backup Options

STEP 1: Choose VMs to Backup

Finding VMs and Containers to Back Up

The left (inventory) pane displays the tree of vCenters and standalone ESXi hosts added to the product (you can refer to the [Configuring Inventory](#) section to learn more). To find items that you want to back up, you can do the following:

- Search for VMs and VMware containers (such as clusters, hosts, folders, and resource pools) by typing a part or the whole item name in the **Search** field:



- Browse through the virtual infrastructure in the left (inventory) pane.

If you cannot find a VM or a container, try the following:

- Make sure the corresponding vCenter or ESX(i) host is added to the inventory as described in the [Adding vCenters and ESX\(i\) hosts](#) section.
- Update the Inventory as described in the [Updating Infrastructure](#) section.


Adding VMs and Containers to the Backup Job

In NAKIVO Backup & Replication, you can add both VMs and VMware containers (such as clusters, hosts, folders, and resource pools) to a backup job. If you add a VMware container to the job:

- All VMs currently available in the selected container will be backed up
- All new VMs that will be created in (or moved to) the container in the future will be automatically added to the job and backed up.

Adding a VMware container to the job enables to ensure that important VMs are always protected.

To add a VM or a container to the job, select the check box next to an inventory item. The VMs and containers you select appear in the right (Selected Items) pane, which displays the item name, path, and processing order:




NOTE: If all disks of a VM are unsupported (such as RDM disks in physical compatibility mode), such a VM will be disabled in the inventory tree and it will not be possible to select such a VM. Refer to the [Supported Hypervisors](#) section to learn more.

Setting the Backup Processing Order

You can specify the order in which VMs should be backed up if the Transporter performing the backup cannot process all VMs of the job simultaneously (either because the Transporter is processing other jobs at the same time or because the number of VM disks in the job exceeds the Transporter's Maximum Load specified during the Transporter creation). By doing so, you can specify to back up the most important VMs first.


To set the VM processing order, simply drag a VM or a container to a new position in the selected VMs pane:



Removing Selected VMs and Containers from the Job

If you want to remove a selected VM or a container from the job, do either of the following:

- In the left pane, deselect the check box next to the item that you would like to remove from the job.
- In the right pane, hover the mouse pointer over the item you want to remove and click the **Remove** button:




Completing the “Choose VMs to Back Up” Step

After you have selected the VMs and containers you want to back up, click **Next** to go to the next step.

STEP 2: Choose a Backup Repository

Setting a Single Backup Repository for All VMs


To back up all VMs (selected in Step 1 of the wizard) to a single backup repository, choose a backup repository from the **Backup repository** dropdown menu:



Setting Different Backup Repositories for VMs

To back up VMs to different backup repositories, follow the steps below:

1. Click **Advanced setup**.
2. In the VM boxes, choose a backup repository individually for each VM:



Mapping Source VMs to Existing Backups

If you have previously backed up a VM and then lost the backup job (due to accidental job deletion or because you need to recreate jobs in a new copy of the product) you can map source VMs to existing backups to avoid running full VM backups again.

To map source VMs to existing backups, follow the steps below:

1. Click **Advanced setup**.
2. From the **Backup repository** dropdown menu, choose a backup repository that contains an existing VM backup.
3. Select the **Use existing backup as a target** option and choose an existing backup from the dropdown menu:

When you run the job, the product will analyze the target VM you have selected, determine how it is different from the source VM, and transfer only the differential data.

NOTE: VM replication mapping can be a time consuming process which can be equal to the time required to create a full VM replication.

After the job completion a new recovery point will be created and existing recovery points will *not* be changed or overwritten.

Excluding VM Disks from the Backup Job

If you do not want to back up some VM disks, you can exclude those disks from the backup job by following the steps below:

1. Click **Advanced setup**.
2. In the VM box, deselect the check box next to the disks you do not want to back up:

NAKIVO
Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

✓ New Backup Job Wizard

STEP 1: Choose VMs to backup

STEP 2: Choose backup repository

Backup repository: Remote repository

To use existing backups as targets, expand the Advanced setup below and configure this per VM.

sy.del Click to collapse

Virtual Machine disks

- ☒ Hard disk 1 (0 KB on datastore1 (7))

Backup repository: Remote repository

☐ Use existing backup as a target

sy.R2stSP1 Click to collapse

Virtual Machine disks

- ☒ Hard disk 1 (17.1 GB on datastore1 (7))
- ☒ Hard disk 2 (1.0 GB on datastore1 (7))
- ☒ Hard disk 3 (1.0 GB on datastore1 (7))
- ☒ Hard disk 4 (32 MB on datastore1 (7))

Backup repository: Remote repository

☐ Use existing backup as a target

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved


Completing the “Choose a Backup Repository” Step

After you have specified the backup repository options, click **Next** to go to the next step.

STEP 3: Specify Backup Schedule

Disabling Scheduled Job Execution

If you only want to start the job manually (without any schedule), select the **Do not schedule, run on demand** check box:




The screenshot shows the 'New Backup Job Wizard' interface for Step 3: Specify backup schedule. The 'Do not schedule, run on demand' checkbox is checked and highlighted with a red circle. The interface includes a header with the NAKIVO logo and navigation links, and a footer with copyright information.

Specifying Job Schedule

Set the backup schedule as follows:

- **To run a job once a day, do the following:**
 - Choose **Run daily/weekly** from the schedule dropdown list:



The screenshot shows the 'New Backup Job Wizard' interface for Step 3: Specify backup schedule. The 'Run daily/weekly' dropdown menu is selected and highlighted with a red circle. The interface includes a header with the NAKIVO logo and navigation links, and a footer with copyright information.

- Specify the time when the job should be started in the **Starting** field.
- Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone which should be used for the job start and end times from the time zone dropdown list.
- Select the days of the week during which the job will be started.
- **To run a job monthly or yearly, do the following:**

- Choose **Monthly/yearly** from the schedule dropdown list:

NAKIVO Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Backup Job Wizard

STEP 1: Choose VMs to backup

STEP 2: Choose backup repository

STEP 3: Specify backup schedule

☐ Do not schedule, run on demand

Run monthly/yearly

Run every **Last** of every **Month**

Starting at: **12:00 AM** Ending: **6:00 AM**

(UTC+02:00, EET) Eastern European Time

March 2014

S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Starts at:

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

- Specify job start schedule in the appropriate fields.
 - Specify the time when the job should be started in the **Starting** field.
 - Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
 - Choose a time zone which should be used for the job start and end times from the time zone dropdown list.
 - Select the days of the week during which the job will be started.
- To run a job multiple times per day, do the following:**
 - Choose **Run periodically** from the schedule dropdown list and then choose a time periods from the appropriate fields:

NAKIVO Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Backup Job Wizard

STEP 1: Choose VMs to backup

STEP 2: Choose backup repository

STEP 3: Specify backup schedule

☐ Do not schedule, run on demand

Run periodically every **30** Minutes

Starting at: **12:00 AM** Ending: **6:00 AM**

(UTC+02:00, EET) Eastern European Time

☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

All days Work days Weekends

March 2014

S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Starts at:

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

- Specify the time when the job should be started in the **Starting** field.

- Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone which should be used for the job start and end times from the time zone dropdown list.

Completing the “Specify backup schedule” Step

After you have specified the schedule options, click **Next** to go to the next step.

STEP 4: Specify Backup Options

Setting Backup Options

Specify backup options for the job as follows:

- **Job name:** Specify a name for the backup job.
- **Enable network acceleration:** If this option is selected, NAKIVO Backup & Replication will use compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to back up over WAN or slow LAN links.
- **Encrypt data during transfer:** If this option is selected, VM data will be protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on machines running Transporters. Select this option if you back up over WAN without a VPN connection.
NOTE: You need at least one Transporter at source and target sites to enable encryption.
- **Run in application-aware mode:** If this option is selected, VM backup will be performed using VMware Guest OS Quiescing (which in turn relies on Microsoft VSS) to ensure that application data is consistent.
- **Enable and use VMware Changed Block Tracking** if this option is selected, NAKIVO Backup & Replication will enable the Changed Block Tracking feature for selected VMs. This feature significantly speeds up the VM backup process.
- **Transport Mode:** Choose between one of the following transport modes for retrieving VM data:
 - **Automatic selection:** If this option is selected, NAKIVO Backup & Replication will try to use direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the product will retrieve data over LAN.
 - **SAN only:** If this option is selected, NAKIVO Backup & Replication will only use the direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the backup job will fail.
 - **LAN only:** If this option is selected, NAKIVO Backup & Replication will only retrieve VM data via LAN.

Setting the Recovery Point Retention Options

After each job run, NAKIVO Backup & Replication creates a recovery point in the backup repository for each VM. A recovery point represents the backed up VM as of a particular moment in time and allows you to recover individual files, application object, or the entire VM from the backup repository.

You can specify how many recovery points should be preserved in the backup repository using Grandfather-Father-Son (GFS) backup rotation scheme. Use the following options:

- **Keep X last recovery points:** Keeps the specified number of last recovery points for each VM in the job.
- **Keep one recovery point per day for X days:** Keeps one last recovery point per day for the specified number of days.
- **Keep one recovery point per week for X weeks:** Keeps the last available backup of every week for the specified number of weeks.
- **Keep one recovery point per month for X months:** Keeps the last available backup of every month for the specified number of months.
- **Keep one recovery point per year for X years:** Keeps the last available backup of every year for the specified number of years.

Truncating Microsoft Exchange Transaction Logs in Backup Job

Microsoft Exchange Server database transaction logs record all changes to an Exchange Server database. Over time, these log files accumulate and can consume all of the available disk space, if not periodically removed. NAKIVO Backup & Replication provides an option to delete (aka truncate) Microsoft Exchange Server logs on the source VMs after job completion.

NOTE: If enabled in a job, this feature will try to truncate Microsoft Exchange logs on all VMs added to the job using the credentials provided. In other words, if you plan to truncate Microsoft Exchange logs on one or more VMs, you need to create a separate job with the VMs, and make sure that all the VMs can be accessed with the credentials you specify in the job.

The transaction logs are deleted after the job completion, so that the log files are available in the VM backup. Note that the product deletes only those transaction logs which are already committed to (available in) the Microsoft Exchange database.

To set up Microsoft Exchange log truncation, do the following:

1. Click **Advanced options**.
2. Select the **Truncate Exchange logs on all VMs** option and then click **settings**:

The screenshot shows the 'STEP 4: Specify backup options' window. Under 'Post job actions', the checkbox 'Truncate Exchange logs on all VMs' is checked, and the 'settings' link next to it is circled in red. Other options include 'Send job reports on each job completion to' (with email address 'notifications@nakivo.com') and 'Transporters' (set to 'Automatically determine source transporter').

3. In the pop-up dialog that appears, do the following:

The 'Log truncation settings' dialog box contains the following fields and options:

- Username: OS username
- Password: OS password
- ☒ Truncate logs on successful backup only
- ☐ Always truncate logs

- a. Specify the username and password in the appropriate fields. These credentials will be used to log in to all source VMs added to the job. The account should have full administrative privileges.
- b. If you want the Microsoft Exchange log files to be always truncated after job completion (regardless of whether the job completes successfully or fails), select the **Always truncate logs** option. Otherwise, leave the **Truncate logs on successful backup only** option selected.
- c. Close the pop-up dialog by clicking anywhere outside the dialog. The entered data will be automatically saved.

Setting up Email Notifications for the Job

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications described in the [Configuring Email Notifications](#) section and provides the ability to configure notifications on the per-job level.

NOTE: To enable this option, configure your Email settings as described in the [Configuring Email Settings](#) section.

To send email notifications, do the following:

1. Click **Advanced options**.

2. Select the **Send job reports on each job completion to** option and specify one or more email addresses in the text field. The semi-colon character should be used to separate multiple email addresses.

STEP 4: Specify backup options

Backup options

Job name:

☐ Enable network acceleration [?](#) ☒ Enable and use VMware CBT [?](#)

☐ Encrypt data during transfer [?](#)

☒ Run in application-aware mode [?](#)

Transport mode: [?](#)

Recovery points retention

☒ Keep last recovery points

☐ Keep one recovery point per day for days

☐ Keep one recovery point per week for weeks

☐ Keep one recovery point per month for months

☐ Keep one recovery point per year for years

[Need more info](#)

Post job actions

☐ Truncate Exchange logs on all VMs [?](#)

☒ Send job reports on each job completion to: [?](#)

Transporters

[?](#)

[Next](#) [Close Wizard](#)

© 2014 NAKIVO, Inc. All rights reserved

Specifying Which Transporters Should be Used by the Backup Job

In NAKIVO Backup & Replication, Transporters perform all of the heavy lifting: read data from the source VM, compress the data, transfer the data over the network, perform data deduplication, and so on. In large and geographically distributed environments multiple Transporters can be deployed to distribute the data protection workload, optimize network traffic, and improve data transfer speeds. Thus, if more than one Transporter is deployed for NAKIVO Backup & Replication, it is important to determine which one should be used to read data from a particular source VM (the target Transporter is always the one that is assigned to the Backup Repository). By default, the product automatically determines which Transporter should be used to read data from the source VM. However, you can manually specify which Transporters should be used for the job. Here's how:

1. Click **Advanced options**.
2. Choose one of the following options:

STEP 4: Specify backup options

Backup options

Job name:

☐ Enable network acceleration [?](#) ☒ Enable and use VMware CBT [?](#)

☐ Encrypt data during transfer [?](#)

☒ Run in application-aware mode [?](#)

Transport mode: [?](#)

Recovery points retention

☒ Keep last recovery points

☐ Keep one recovery point per day for days

☐ Keep one recovery point per week for weeks

☐ Keep one recovery point per month for months

☐ Keep one recovery point per year for years

[Need more info](#)

Post job actions

☐ Truncate Exchange logs on all VMs [?](#)

☒ Send job reports on each job completion to: [?](#)

Transporters

[?](#)

[Next](#) [Close Wizard](#)

© 2014 NAKIVO, Inc. All rights reserved

- **Automatically determine source transporter:** The product will automatically determine which Transporters are the closest to source hosts (the hosts that run selected VMs) and will use those Transporters to retrieve data from source VMs.
- **Set a source transporter for all VMs:** Select this option to manually specify a single Transporter that will be used to retrieve data from source VMs.
- **Manually set transporters for source hosts:** Select this option to manually specify which Transporter should be used to retrieve data from each source host.


NOTE: *The target Transporter for the backup job will always be the Transporter assigned to the backup repository.*

Completing the New Backup Job wizard

Click **Finish** to complete the job creation.

Viewing Backup Job Information

Job information is aggregated in job dashboards. To open a dashboard for a particular job, click the job in the navigation pane:




Backup Job Widgets

Job details are available in job widgets. Refer to the sections below to learn about backup job widgets.

The “Backup Job Status” Widget

The “Backup Job Status” widget displays the current status of the job and the status of the last job run:



The “Source VMs” Widget

The “Source VMs” widget displays the list of VMs added to the backup job:

Source VMs		▼ Less details
2003forAmazonTest	Runs on demand	
Host:	192.168.77.213 > Ninja-Site1 > 10.30.28.10	
Datastores:	esx10datastore	
Networks:	192.168.77.0	
testSY_win7	Runs on demand	
Host:	192.168.77.213 > Ninja-Site1 > 10.30.28.10	
Datastores:	esx10datastore	
Networks:	192.168.77.0	

The following data is displayed for each VM:

- Power status of the source VMs
- Path to the host where each source VM is located
- Datastores where the source VM files are located
- Networks to which each source VM is connected

The “Backups” Widget

The “Backups” widget displays the list of backups which are created (or will be created) by the job:

Backups		▼ Less details
 2003forAmazonTest		
Repository:	Onboard repository - 150 GB (96.7% free)	
Recovery points:	2 (last on Jul 12 at 12:07 PM UTC +03:00)	
Size after recovery:	12 GB	
 testSY_win7		Not created yet
Repository:	Onboard repository - 150 GB (96.7% free)	
Recovery points:	Not created yet	
Size after recovery:	17 GB	

The following data is displayed:

- Names of backups
- Repository where a backup is located
- Number of recovery points for a backup
- Size of the recovered VM

The “Events” Widget

The “Events” widget displays all job-related events:

Events	
The job has finished The "New backup job" job has successfully finished.	Jul 12 at 12:07 PM
VM backup has started The "New backup job" job has started backing up the "2003forAmazonTest" VM.	Jul 12 at 12:07 PM
The job has started The "New backup job" job is starting its adhoc run. The job has run 1 time(s) previously.	Jul 12 at 12:07 PM
The job is running The "New backup job" job is currently running.	Jul 12 at 12:07 PM

The “Backup Job Stats” Widget

The “Backup Job Stats” widget displays statistical information about the job:


Job Stats	
2 source VMs	
4 disks	
21 GB - size of source VMs	
0.0% - data transfer compression on last run	
0 minutes - duration of last run	
0 minutes - average duration	

The following data is displayed:

- The number of source VMs backed up by the job
- The number of disks backed up by the job
- Compression ratio achieved while transferring data over the network
- The duration of last run
- The average duration of job runs


The “Schedule and Options” Widget

The “Schedule and Options” widget displays the job schedule and backup options you have specified during the job creation:



The “Backup Repositories” Widget

The “Backup Repositories” widget displays all backup repositories to which the job backs up source VMs:




For each backup repository, the following data is displayed:

- The amount of free space on the backup repository
- The amount of stored data
- Data compression ratio
- Data deduplication ratio

The “Average Bandwidth Utilization” Widget


The “Average Bandwidth Utilization” widget graphically displays the average data transfer speed for the last 10 job runs. The widget calculates and shows the average speed at which the data has been transferred. The information is based on the amount of time during which the job has completed and the amount of data transferred.



If one of the runs has failed, the widget displays data for that run as a red bar. If one of the runs has been stopped, the widget displays data for that run as a yellow bar.

The “Transferred Data” Widget


The “Transferred Data” widget graphically displays the amount of transferred data for the last 10 job runs.



If one of the runs has failed, the widget displays data for that run as a red bar. If one of the runs has been stopped, the widget displays data for that run as a yellow bar.


The “Source Transporters” Widget

The source Transporters widget lists all Transporters that retrieve data from the source VMs and sends the data to the target Transporters:



The “Target Transporters” Widget

The target Transporters widget lists all Transporters that receive data from source Transporters and writes the data to the backup repository:




Managing VM Backups

Deleting Recovery Points

In NAKIVO Backup & Replication, you can permanently delete a recovery point in a VM backup. Deleting a recovery point does not affect any other VM backup recovery points.

To permanently delete a recovery point, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab and click a backup repository.
3. Click the backup where you want to delete a recovery point.
4. Hover the mouse pointer over the recovery point you want to delete and click **Delete**:



NOTE: A recovery point cannot be deleted while the source VM backup job is running.

5. Click **Delete Recovery Point** in the message that appears.


NOTE: The space occupied by the recovery point will be marked as “free” and reused by new data blocks on the next job runs, though the actual size of the backup repository may not change. To free up the space occupied by the recovery point, you can reclaim free space as described in the [Verifying VM Backups](#). The backup verification process reads each block of data inside the backup repository and ensures that it is identical to the data block that was read on the source VM during the backup. This way the product verifies that files, application objects, and entire VMs can be recovered from all VM backups.

NOTE: Before the backup verification begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the backup verification process is in progress.

To verify all VM backups in a repository, follow the steps below:


6. Click **Configuration** in the upper right corner of the product.
7. Go to the **Backup repositories** tab and click a backup repository.

8. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The backup verification process cannot be started if a job which backs up to this backup repository is running.

In the dialog box that appears, select the **Perform full data verification** option and click **Start**:



The backup verification process is started.

Dealing with Corrupted Recovery Points

If the backup verification process finds that a recovery point is corrupted, do the following:


- If the last recovery point is corrupted, follow the steps below to repair it:
 - d. Delete the backup job (as described in the [Deleting Jobs](#) section) or edit the backup job (as described in the [Editing Jobs](#) section) and remove the source VM that has the corrupted recovery point.
 - e. Create a new backup job and map the source VM to the backup with the corrupted recovery point as described in the [Mapping Source VMs to Existing Backups](#) section.
 - f. Run the job. The recovery point will be repaired using the data from the source VM.
- If the corrupted recovery point is not the last one, it cannot be repaired. Delete such recovery point.

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the backup verification process, follow the steps below:

9. Click **Configuration** in the upper right corner of the product.
10. Go to the Backup **repositories** tab and click a backup repository.
11. In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The backup verification process is stopped.

Reclaiming Backup Repository Space section.


Deleting VM Backups

In NAKIVO Backup & Replication, you can permanently delete a VM backup (including all of its recovery points) if it is available in a backup repository.

NOTE: You can only delete a VM backup if you have deleted the corresponding backup job or edited the backup job so it does not include the backup's source VM.

To permanently delete a VM backup, follow the steps below:

1. Click **Configuration** in the upper right corner of the product.
2. Go to the **Backup repositories** tab and click a backup repository.
3. Click the backup you want to delete.
4. In the backup title, click **Delete**:




5. Click **Delete Backup** in the message that appears.

NOTE: The space occupied by the backup will be marked as “free” and reused by new data blocks on the next job runs, though the actual size of the backup repository may not change. To free up the space occupied by the backup, you can reclaim free space as described in the [Verifying VM Backups](#). The backup verification process reads each block of data inside the backup repository and ensures that it is identical to the data block that was read on the source VM during the backup. This way the product verifies that files, application objects, and entire VMs can be recovered from all VM backups.

NOTE: Before the backup verification begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the backup verification process is in progress.


To verify all VM backups in a repository, follow the steps below:

6. Click **Configuration** in the upper right corner of the product.
7. Go to the **Backup repositories** tab and click a backup repository.
8. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The backup verification process cannot be started if a job which backs up to this backup repository is running.

In the dialog box that appears, select the **Perform full data verification** option and click **Start**:



The backup verification process is started.

Dealing with Corrupted Recovery Points

If the backup verification process finds that a recovery point is corrupted, do the following:


- If the last recovery point is corrupted, follow the steps below to repair it:
 - g. Delete the backup job (as described in the [Deleting Jobs](#) section) or edit the backup job (as described in the [Editing Jobs](#) section) and remove the source VM that has the corrupted recovery point.
 - h. Create a new backup job and map the source VM to the backup with the corrupted recovery point as described in the [Mapping Source VMs to Existing Backups](#) section.
 - i. Run the job. The recovery point will be repaired using the data from the source VM.
- If the corrupted recovery point is not the last one, it cannot be repaired. Delete such recovery point.

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the backup verification process, follow the steps below:

9. Click **Configuration** in the upper right corner of the product.
10. Go to the Backup **repositories** tab and click a backup repository.
11. In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The backup verification process is stopped.

Reclaiming Backup Repository Space section.

VM Backup How-Tos

How to Stage (Seed) Initial VM Backup

Because VMs are often large, the initial VM backup can be slow and time consuming when done over WAN, and can also put undesirable load on the network when done over LAN. To speed up the initial backup and save network bandwidth, you can perform staged backup: run the initial backup on a removable media (such as an external USB hard drive), transfer the media to a new location, and then run incremental backup to the new location.

To perform staged backup, follow the steps below:

1. Create a new backup repository as described in the [Adding Backup Repositories](#) section.
2. Create a new backup job with the following options:
 - a. In [STEP 1: Choose VMs to Backup](#), choose VMs for which you want to run staged backup.
 - b. In [STEP 2: Choose a Backup Repository](#), choose the backup repository you have created.
 - c. In [STEP 3: Specify Backup Schedule](#), select the **Do not schedule, run on demand** option.
 - d. In [STEP 4: Specify Backup Options](#), specify required backup options.
3. Run the backup job as described in the [Running Jobs on Demand](#) section.
4. After the job has completed, move the backup repository to a new location as described in the [How to Move a Backup Repository to a New Location](#) section.
5. If required, edit the backup job as described in the [Editing Jobs](#) section and specify a schedule for the backup job in [STEP 3: Specify Backup Schedule](#).

How to Back Up to Two Removable Disks in Rotation

One of the data protection best practices suggests to always keep a copy of important backups offsite. To achieve that, one of the approaches is to back up VMs to two removable disks in rotation, and keep one of them outside the infrastructure. Here's how you can achieve this with NAKIVO Backup & Replication:

Preparing Backup Repositories

To back up to two disks in rotation, you need to create a backup repository on each disk.

NOTE: *If you create just one backup repository on Disk 1 and copy/paste it to Disk 2, you will not be able to back up to Disk 2 in rotation.*

1. Prepare two disks to which you plan to backup VMs.
2. Connect or mount Disk 1 to a machine where a Transporter is installed.
3. Create a *backup repository 1* on Disk 1 as described in the [Adding a Local Folder as a Backup Repository](#) section.
4. Connect or mount Disk 2 to a machine where a Transporter is installed.
5. Create a *backup repository 2* on Disk 2 as described in the [Adding a Local Folder as a Backup Repository](#) section.

Preparing Backup Jobs

To back up to two disks in rotation, you need to create two separate backup jobs: a *backup job 1* that backs up your VMs to Disk 1 and a *backup job 2* that backs up the same VMs to Disk 2. Here's how:

1. Create a *backup job 1* for Disk 1:
 - a. Make sure that Disk 1 is connected or mounted to the machine where the Transporter is installed.
 - b. If the backup repository 1 was detached, attach it as described in the [Attaching Backup Repositories](#) section.
 - c. Start the New Backup Job Creation Wizard as described in the [Creating Backup Jobs](#) section.
 - d. In [STEP 1: Choose VMs to Backup](#), choose VMs that you want to back up.
 - e. In [STEP 2: Choose a Backup Repository](#), choose the *backup repository 1* you have created.
 - f. In [STEP 3: Specify Backup Schedule](#), select the **Do not schedule, run on demand** option.
 - g. In [STEP 4: Specify Backup Options](#), specify required backup options.
2. Create a *backup job 2* for Disk 2:
 - a. Make sure that Disk 2 is connected or mounted to the machine where the Transporter is installed.
 - b. If the backup repository 2 was detached, attach it as described in the [Attaching Backup Repositories](#) section.
 - c. Start the New Backup Job Creation Wizard as described in the [Creating Backup Jobs](#) section.
 - d. In [STEP 1: Choose VMs to Backup](#), choose VMs that you want to back up.
 - e. In [STEP 2: Choose a Backup Repository](#), choose the *backup repository 2* you have created.
 - f. In [STEP 3: Specify Backup Schedule](#), select the **Do not schedule, run on demand** option.
 - g. In [STEP 4: Specify Backup Options](#), specify required backup options.

Backing up to Removable Disks in Rotation


To back up to a removable disk, do the following:

1. Connect or mount Disk 1 to the machine where the Transporter is installed.
2. Attach the *backup repository 1* as described in the [Attaching Backup Repositories](#) section.
3. Run the *backup job 1*.
4. After the backup job 1 has completed, detach the backup repository 1 as described in the [Detaching Backup Repositories](#) section.
5. Disconnect or unmount the Disk 1 – it's now ready to be relocated offsite.
6. Repeat the above steps for Disk 2.

How to Improve Backup Performance Over WAN

To achieve maximum performance over WAN, follow the steps below:

1. Install a new Transporter instance in the remote site where the backup repository is located.
2. Add the newly installed Transporter to the product as described in the [Adding Installed Transporters](#) section.
3. Edit the backup repository as described in the [Editing Backup Repositories](#) section and choose the newly added Transporter from the **Assigned Transporter** dropdown menu. At this point, the VM backup data flow is as follows:



4. Edit the backup job as described the [Editing Jobs](#) section and make sure the **Enable network acceleration** option is selected in [STEP 4: Specify Backup Options](#).

How to Stop Backing Up a VM

If you no longer want to back up a VM, you can remove the VM from the backup job by following the steps below:

1. Edit the backup job as described in the [Editing Jobs](#) section.
2. Remove the VM from the [STEP 1: Choose VMs to Backup](#), as described in the [Removing Selected VMs and Containers from the Job](#) section.

Troubleshooting VM Backup

Reducing the Number of Recovery Points Does Not Free Up Space


Refer to the [Verifying VM Backups](#)

The backup verification process reads each block of data inside the backup repository and ensures that it is identical to the data block that was read on the source VM during the backup. This way the product verifies that files, application objects, and entire VMs can be recovered from all VM backups.

NOTE: Before the backup verification begins, the backup repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such backup repository will fail while the backup verification process is in progress.


To verify all VM backups in a repository, follow the steps below:

3. Click **Configuration** in the upper right corner of the product.
4. Go to the Backup **repositories** tab and click a backup repository.
5. In the title of the backup repository, click **Manage** and then click **Start Maintenance**:



NOTE: The backup verification process cannot be started if a job which backs up to this backup repository is running.

In the dialog box that appears, select the **Perform full data verification** option and click **Start**:



The backup verification process is started.

Dealing with Corrupted Recovery Points

If the backup verification process finds that a recovery point is corrupted, do the following:

- If the last recovery point is corrupted, follow the steps below to repair it:


- h. Delete the backup job (as described in the [Deleting Jobs](#) section) or edit the backup job (as described in the [Editing Jobs](#) section) and remove the source VM that has the corrupted recovery point.
 - i. Create a new backup job and map the source VM to the backup with the corrupted recovery point as described in the [Mapping Source VMs to Existing Backups](#) section.
 - j. Run the job. The recovery point will be repaired using the data from the source VM.
- If the corrupted recovery point is not the last one, it cannot be repaired. Delete such recovery point.

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example to run a recovery job, move your backup repository to a new location, or put your backup storage on maintenance).

To stop the backup verification process, follow the steps below:

6. Click **Configuration** in the upper right corner of the product.
7. Go to the Backup **repositories** tab and click a backup repository.
8. In the title of the backup repository, click **Manage** and then click **Stop maintenance**:



The backup verification process is stopped.

Reclaiming Backup Repository Space section for more information on this issue.

VM Cannot be Selected for Backup

VMs that have unsupported disk types are disabled and cannot be selected for backup. Refer to the VM Backup and Replication Requirements section for more information.

Quiesced Snapshot Creation Failure

If a quiesced snapshot fails to create:

- Update your ESX(i) hosts to the latest version
- Update VMware tools to the latest version
- Update applications and databases running on the VM to their latest versions

Changed Block Tracking Failure

Changed Block Tracking failure can occur if your infrastructure is not updated to the latest version. To resolve this issue:

- Update your ESX(i) hosts to the latest version
- Update VMware tools to the latest version

A Problem with a Backup Infrastructure

Network connectivity disruptions, change of credentials and permissions, improper backup repository detach, and even third-party software can cause metadata and backup data to be locked, partially written, or even corrupted in a backup repository. To ensure data integrity at all times, the product periodically runs quick self-checks on backup repositories.

In case a quick self-check finds symptoms of a problem in a backup infrastructure – such as incorrect timestamps on metadata and data files – the backup repository is locked to prevent any potential data corruption. To unlock the backup repository, run the full self-check as described in the [Running Backup Repository Self-Healing](#) section. If this problem persists, please refer to NAKIVO Technical Support for further investigation.

Replication

This section covers the following topics:

- [Staging VM Replication](#)
- [Creating Replication Jobs](#)
- [Viewing Replication Job Information](#)

Staging VM Replication

Because VMs are often large, the initial (full) VM replication can be slow, time consuming, and can put undesirable load on the network. To speed up the initial VM replication and save network bandwidth, you can perform staged replication: transfer VMs to your target site using a removable media (such as an external USB hard drive) and then create a new replication job that will use the transferred VMs as a target and perform only incremental replication.


To stage VM replication, follow the steps below:

1. Put VMs that you want to replicate to a removable media (such as an external USB hard drive) using backup, replication, or any other means.
2. Transfer the media to the target location.
3. Add (recover) the VMs to the desired server and datastore.
4. Create a new replication job and map the source VMs to the transferred VMs as described in the [Mapping Source VMs to Existing Replicas](#)

If you want to perform staged replication as described in the [Staging VM Replication](#) section or if you have lost the replication job (due to accidental job deletion or because you need to recreate jobs in a new copy of the product) you can map source VMs to existing replicas to avoid running full VM replication again.

To map source VMs to existing VMs, follow the steps below:

1. Click **Advanced setup**.
2. Open the **Choose a target logical unit** dropdown menu and select the **Use existing VM as a target** option.



3. The **Choose a target logical unit** dropdown menu contents is updated to include VMs.
4. Select the VM that should be used as a target for replication.

When you run the job, the product will analyze the target VM you have selected, determine how it is different from the source VM, and transfer only the data differential data.

NOTE: VM replication mapping can be a time consuming process which can be equal to the time required to create a full VM replication.

After the job completion a new recovery point will be created and existing recovery points will *not* be changed or overwritten.

Excluding a VM Disks from the Replication Job

If you do not want to replicate some disks of a VM, you can exclude those disks from the replication job. Follow the steps below:

1. Click **Advanced setup**.
2. In the VM box, expand the **Disks** box.
3. Click the drop-down menu next to the disk that you want to skip.
4. From the dropdown, choose the **Skip this disk** option:

NAKIVO
Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Replication Job Wizard

STEP 1: Choose VMs to replicate

STEP 2: Choose replica location

Target: 10.30.28.23 | Datastore: esx21datastore | Network: 10.30.29.0

To use existing VMs as targets, expand the Advanced setup below and configure this per VM. Target is selected in Choose a target logical unit menu.

sy.R2stSP1	10.30.28.23
VM: 2 CPU, 4.0 GB RAM	Host: 1 CPU, 8.0 GB RAM
Path: 10.30.29.12 > QA datacenter > 10.30.29.46	Path: 10.30.29.12 > IT
▼ 4 disks: 19.2 GB	▼ 4 disks: esx21datastore (432.2 GB free)
Hard disk 1: 17.1 GB on datastore1 (7)	Hard disk 1: esx21datastore (432.2 GB free)
Hard disk 2: 1.0 GB on datastore1 (7)	Hard disk 2: esx21datastore (432.2 GB free)
Hard disk 3: 1.0 GB on datastore1 (7)	Hard disk 3: esx21datastore (432.2 GB free)
Hard disk 4: 32 MB on datastore1 (7)	Hard disk 4: esx21datastore (432.2 GB free)
VM file: datastore1 (7)	VM file: esx21datastore (432.2 GB free)
▶ 1 NIC	▶ 1 NIC: 10.30.29.0


Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

5. section.

Creating Replication Jobs

In NAKIVO Backup & Replication you can replicate VMs by creating a job that specifies which VMs should be replicated, where the replicas should be located, how often the replication should occur, and what replication options should be used. To create a replication job, click **Create** and then click **Replication Job** on any dashboard:




The New Replication Job Wizard opens. Complete the wizard as described below to create a replication job.

STEP 1: Choose VMs to Replicate

Finding VMs and Containers to Replicate

The left (inventory) pane displays the tree of vCenters and standalone ESXi hosts added to the product (you can refer to the [Configuring Inventory](#) section to learn more). To find items that you want to replicate you can do the following:

- Search for VMs and VMware containers (such as clusters, hosts, folders, and resource pools) by typing a part or the whole item name in the **Search** field:



- Browse the Inventory pane to locate an item.

If you cannot find a VM or a container:

- Make sure the corresponding vCenter or ESX(i) host is added to the inventory as described in the [Adding vCenters and ESX\(i\) hosts](#) section.
- Update the Inventory as described in the [Updating Infrastructure](#) section.


Adding VMs and Containers to the Replication Job

In NAKIVO Backup & Replication, you can add both VMs and VMware containers (such as clusters, hosts, folders, and resource pools) to a replication job. If you add a VMware container to the job:

- All VMs currently available in the selected container will be replicated
- All new VMs that will be created in (or moved to) the container in the future will be automatically added to the job and replicated.

Adding a VMware container to the job enables to ensure that important VMs are always protected.

To add a VM or a container to the job, select the check box next to an inventory item. The VMs and containers you select appear in the right (Selected Items) pane, which displays the item name, path, and processing order:




NOTE: If all disks of a VM are unsupported (such as RDM disks in physical compatibility mode), this VM will be disabled in the inventory tree and will not be possible to select. Refer to the [VM Backup and Replication Requirements](#) section to learn more.

The VMs and containers you select appear in the right (Selected VMs) pane.

Setting the Replication Processing Order

You can specify the order in which VMs should be replicated if the Transporter performing replication cannot process all VMs of the job simultaneously (either because the Transporter is processing other jobs at the same time or because the job contains more VM disks than a Transporter's Maximum Load specified during the Transporter creation). By setting the replication processing order, you can specify to replicate the most important VMs first.


To set the processing order, simply drag a VM or a container to a new position in the selected VMs pane:



Removing Selected VMs and Containers from the Job

If you want to remove a selected VM or a container from the job, do either of the following:

- In the left pane, deselect the check box next to the item you would like to remove from the job.
- In the right pane, hover the mouse pointer over the item you want to remove and click the **Remove** button:



Completing the “Choose VMs to Replicate” Step


After you have selected the VMs you want to replicate, click **Next** to go to the next step.

STEP 2: Choose Replica Location

Setting the Same Host, Datastore, and Network for All Replicas

To replicate all VMs to the same container and datastore as well as connect all replicas to the same networks, follow the steps below:


1. Choose a cluster, host, or resource pool from the **Target** dropdown menu.
2. Choose a datastore from the **Datastore** dropdown menu.
3. Choose a network from the **Network** dropdown menu.



Setting Different Options for Replicas

To specify different replication options for VMs, follow the steps below:

1. Click **Advanced setup**.
2. Choose a target host, target datastore, and network for each VM:




Mapping Source VMs to Existing Replicas

If you want to perform staged replication as described in the [Staging VM Replication](#) section or if you have lost the replication job (due to accidental job deletion or because you need to recreate jobs in a new copy of the product) you can map source VMs to existing replicas to avoid running full VM replication again.

To map source VMs to existing VMs, follow the steps below:

3. Click **Advanced setup**.
4. Open the **Choose a target logical unit** dropdown menu and select the **Use existing VM as a target** option.



5. The **Choose a target logical unit** dropdown menu contents is updated to include VMs.
6. Select the VM that should be used as a target for replication.

When you run the job, the product will analyze the target VM you have selected, determine how it is different from the source VM, and transfer only the data differential data.

NOTE: VM replication mapping can be a time consuming process which can be equal to the time required to create a full VM replication.

After the job completion a new recovery point will be created and existing recovery points will *not* be changed or overwritten.

Excluding a VM Disks from the Replication Job

If you do not want to replicate some disks of a VM, you can exclude those disks from the replication job. Follow the steps below:

7. Click **Advanced setup**.
8. In the VM box, expand the **Disks** box.
9. Click the drop-down menu next to the disk that you want to skip.
10. From the dropdown, choose the **Skip this disk** option:

NAKIVO

Backup & Replication

[Welcome, admin](#) |
 [Configuration](#) |
 [Support](#) |
 [Help](#) |
 [Log out](#)

New Replication Job Wizard

STEP 1: Choose VMs to replicate

STEP 2: Choose replica location

Target

10.30.28.23

Datastore

esx21datastore

Network

10.30.29.0

To use existing VMs as targets, expand the Advanced setup below and configure this per VM. Target is selected in Choose a target logical unit menu.

sy.R2stSP1 VM: 2 CPU, 4.0 GB RAM Path: 10.30.29.12 > QA datacenter > 10.30.29.46 ▼ 4 disks: 19.2 GB Hard disk 1: 17.1 GB on datastore1 (7) Hard disk 2: 1.0 GB on datastore1 (7) Hard disk 3: 1.0 GB on datastore1 (7) Hard disk 4: 32 MB on datastore1 (7) VM file: datastore1 (7) ► 1 NIC	10.30.28.23 ▼ Host: 1 CPU, 8.0 GB RAM Path: 10.30.29.12 > IT ▼ 4 disks: esx21datastore (432.2 GB free) Hard disk 1: esx21datastore (432.2 GB free) Hard disk 2: esx21datastore (432.2 GB free) Hard disk 3: esx21datastore (432.2 GB free) Hard disk 4: esx21datastore (432.2 GB free) VM file: esx21datastore (432.2 GB free) ► 1 NIC: 10.30.29.0
---	---

[Collapse](#)

Next

Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

Completing the “Choose replica location” Step

After you have specified replica location options, click **Next** to go to the next step.

STEP 3: Specify Replication Schedule

Disabling Scheduled Job Execution

If you only want to start the job manually (without any schedule), select the **Do not schedule, run on demand** check box:

NAKIVO Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Replication Job Wizard

STEP 1: Choose VMs to replicate

STEP 2: Choose replica location

STEP 3: Specify replication schedule

☒ Do not schedule, run on demand

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

Specifying Job Schedule

Set the backup schedule as follows:

- **To run a job once a day, do the following:**
 - Choose **Run daily/weekly** from the schedule dropdown list:

NAKIVO Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Replication Job Wizard

STEP 1: Choose VMs to replicate

STEP 2: Choose replica location

STEP 3: Specify replication schedule

☐ Do not schedule, run on demand

☒ Run daily/weekly

Starting at: 12:00 AM Ending: 6:00 AM

(UTC+02:00, EET) Eastern European Time

☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

All days Work days Weekends

March 2014

S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Starts at:

AM

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

- Specify the time when the job should be started in the **Starting** field.
- Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone which should be used for the job start and end times from the time zone dropdown list.
- Select the days of the week during which the job will be started.
- **To run a job monthly or yearly, do the following:**

- Choose **Monthly/yearly** from the schedule dropdown list:

NAKIVO
Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Replication Job Wizard

STEP 1: Choose VMs to replicate

STEP 2: Choose replica location

STEP 3: Specify replication schedule

☐ Do not schedule, run on demand

Run monthly/yearly (highlighted)

Run every **Last** of every **Month**

Starting at: **12:00** **AM** Ending: **6:00** **AM**

(UTC+02:00, EET) Eastern European Time

March 2014 calendar: 4, 28 highlighted

Starts at: (clock icon)

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

- Specify job start schedule in the appropriate fields.
 - Specify the time when the job should be started in the **Starting** field.
 - Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
 - Choose a time zone which should be used for the job start and end times from the time zone dropdown list.
 - Select the days of the week during which the job will be started.
- To run a job multiple times per day, do the following:**
 - Choose **Run periodically** from the schedule dropdown list and then choose a time periods from the appropriate fields:

NAKIVO
Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Replication Job Wizard

STEP 1: Choose VMs to replicate

STEP 2: Choose replica location

STEP 3: Specify replication schedule

☐ Do not schedule, run on demand

Run periodically (highlighted) every **30** **Minutes**

Starting at: **12:00** **AM** Ending: **6:00** **AM**

(UTC+02:00, EET) Eastern European Time

☐ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☐ Sat

All days Work days Weekends

March 2014 calendar: 4, 5, 6, 7, 10, 11, 12, 13, 14, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31 highlighted

Starts at: (clock icon)

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

- Specify the time when the job should be started in the **Starting** field.

- Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone which should be used for the job start and end times from the time zone dropdown list.
- Select the days of the week during which the job will be started.

Completing the “Choose Replication Schedule” Step

After you have specified schedule options, click **Next** to go to the next step.

STEP 4: Specify Replication Options

Setting Replication Options

Specify replication options for the job as follows:

- **Job name:** Specify a name for the replication job.
- **Enable network acceleration:** If this option is selected, NAKIVO Backup & Replication will use compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to replicate over WAN or slow LAN links.
- **Encrypt data during transfer:** If this option is selected, VM data will be protected with AES 256 encryption while traveling over the network. Data encryption increases the replication time and CPU load on machines running Transporters. Select this option if you replicate over WAN without a VPN connection.

NOTE: You need at least one Transporter at source and target sites to enable encryption.

- **Run in application-aware mode:** If this option is selected, VM replication will be performed using VMware Guest OS Quiescing (which in turn relies on Microsoft VSS) to ensure that application data is consistent.
- **Create only thin disks for replicas:** If this option is selected, only thin disks will be created on replicas, regardless of the disk types of the original source VM.
- **Enable and use VMware Changed Block Tracking** if this option is selected, NAKIVO Backup & Replication will enable the Changed Block Tracking feature for selected VMs. This feature significantly speeds up VM replication.
- **Transport Mode:** Choose between one of the following transport modes for retrieving VM data:
 - **Automatic selection:** If this option is selected, NAKIVO Backup & Replication will try to use direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the product will retrieve data over LAN.
 - **SAN only:** If this option is selected, NAKIVO Backup & Replication will only use the direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the replication job will fail.
 - **LAN only:** If this option is selected, NAKIVO Backup & Replication will only retrieve VM data via LAN.

Setting the Recovery Point Retention Options

After each job run, NAKIVO Backup & Replication can create a recovery point (snapshot) on the replica VM. You can specify how many recovery points should be preserved in the backup repository using Grandfather-Father-Son (GFS) backup rotation scheme.

NOTE: Up to 30 recovery points in total can be created on a Replica VM.

Use the following options to specify a retention policy:

The screenshot shows the 'New Replication Job Wizard' in the NAKIVO Backup & Replication interface. The wizard is at Step 4: Specify replication options. The 'Recovery points retention' section is highlighted with a red box. It contains several options: 'Keep 10 last recovery points' (checked), 'Keep one recovery point per day for 10 days' (unchecked), 'Keep one recovery point per week for 4 weeks' (unchecked), 'Keep one recovery point per month for 12 months' (unchecked), and 'Keep one recovery point per year for 3 years' (unchecked). Other options include 'Enable network acceleration', 'Create only thin disks on replicas', 'Encrypt data during transfer', 'Enable and use VMware CBT', and 'Run in application-aware mode'. The 'Transport mode' is set to 'Automatic selection'. There are 'Next' and 'Close Wizard' buttons at the bottom right.

- **Keep X last recovery points:** Keeps the specified number of last recovery points for each VM in the job.
- **Keep one recovery point per day for X days:** Keeps one last recovery point per day for the specified number of days.
- **Keep one recovery point per week for X weeks:** Keeps the last available backup of every week for the specified number of weeks.
- **Keep one recovery point per month for X months:** Keeps the last available backup of every month for the specified number of months.
- **Keep one recovery point per year for X years:** Keeps the last available backup of every year for the specified number of years.

Specifying VM Replica Names

NAKIVO Backup & Replication provides the ability to change VM Replica names so that the VM Replicas can be easily distinguished from the source VMs. By default, the text “- replica” is appended to the end of the VM Replica name.

To change VM Replica names, do the following:

1. Click **Advanced options**.
2. Choose between the following options:

STEP 4: Specify replication options

Replication options

Job name:

☐ Enable network acceleration ?
 ☐ Create only thin disks on replicas ?
 ☐ Encrypt data during transfer ?
 ☒ Enable and use VMware CBT ?
 ☒ Run in application-aware mode ?

Transport mode: ?

Recovery points retention

☒ Keep last recovery points
 ☐ Keep one recovery point per day for days
 ☐ Keep one recovery point per week for weeks
 ☐ Keep one recovery point per month for months
 ☐ Keep one recovery point per year for years

[Need more info](#)

Replica names

☒ Append after
☐ Leave replica names as is
 ☐ Enter custom replica names

Post job actions

☐ Truncate Exchange logs on all VMs ?
 ☐ Send job reports on each job completion to: ?

Transporters

?

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

- **Append X before/after VM name** – source VM names will be used for replica names and the specified text will be added before or after the replica name.
- **Leave replica names as is** – replica names will be identical to the source VM names.
- **Enter custom replica names** – provides the ability to enter custom names for replicas.

Truncating Microsoft Exchange Transaction Logs in Replication Job

Microsoft Exchange Server database transaction logs record all changes to an Exchange Server database. Over time, these log files accumulate and can consume all of the available disk space, if not periodically removed. NAKIVO Backup & Replication provides an option to delete (aka truncate) Microsoft Exchange Server logs on the source VMs after job completion.

NOTE: If enabled in a job, this feature will try to truncate Microsoft Exchange logs on all VMs added to the job using the credentials provided. In other words, if you plan to truncate Microsoft Exchange logs on one or more VMs, you need to create a separate job with the VMs, and make sure that all the VMs can be accessed with the credentials you specify in the job.

The transaction logs are deleted after the job completion, so that the log files are available in the VM Replica. Note that the product deletes only those transaction logs which are already committed to (available in) the Microsoft Exchange database.

To set up Microsoft Exchange log truncation, do the following:

1. Click **Advanced options**.
2. Select the **Truncate Exchange logs on all VMs** option and then click **settings**:

STEP 4: Specify replication options

Replication options

Job name:

☐ Enable network acceleration ?
 ☐ Create only thin disks on replicas ?
 ☐ Encrypt data during transfer ?
 ☒ Enable and use VMware CBT ?
 ☒ Run in application-aware mode ?

Transport mode: ?

Recovery points retention

☒ Keep last recovery points
☐ Keep one recovery point per day for days
☐ Keep one recovery point per week for weeks
☐ Keep one recovery point per month for months
☐ Keep one recovery point per year for years

[Need more info](#)

Replica names

☒ Append after
☐ Leave replica names as is
☐ Enter custom replica names

Post job actions

☒ Truncate Exchange logs on all VMs ?
☐ Send job reports on each job completion to: ?

Transporters

?

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

3. In the pop-up dialog that appears, do the following:

Log truncation settings

Username:

Password:

☒ Truncate logs on successful backup only
☐ Always truncate logs

- k. Specify the username and password in the appropriate fields. These credentials will be used to log in to all source VMs added to the job. The account should have full administrative privileges.
- l. If you want the Microsoft Exchange log files to be always truncated after job completion (regardless of whether the job completes successfully or fails), select the **Always truncate logs** option. Otherwise, leave the **Truncate logs on successful backup only** option selected.
- m. Close the pop-up dialog by clicking anywhere outside the dialog. The entered data will be automatically saved.

Setting up Email Notifications for the Job

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications described in the [Configuring Email Notifications](#) section and provides the ability to configure notifications on the per-job level.

NOTE: To enable this option, configure your Email settings as described in the [Configuring Email Settings](#) section.

To send email notifications, do the following:

1. Click **Advanced options**.
2. Select the **Send job reports on each job completion to** option and specify one or more email addresses in the text field. The semi-colon character should be used to separate multiple email addresses.

STEP 4: Specify replication options

Replication options

Job name:

☐ Enable network acceleration [?](#)
☐ Create only thin disks on replicas [?](#)

☐ Encrypt data during transfer [?](#)
☒ Enable and use VMware CBT [?](#)

☒ Run in application-aware mode [?](#)

Transport mode: [?](#)

[Need more info](#)

Recovery points retention

☒ Keep last recovery points

☐ Keep one recovery point per day for days

☐ Keep one recovery point per week for weeks

☐ Keep one recovery point per month for months

☐ Keep one recovery point per year for years

Replica names

☒ Append after

☐ Leave replica names as is

☐ Enter custom replica names

Post job actions

☐ Truncate Exchange logs on all VMs [?](#)

☒ Send job reports on each job completion to:

Transporters

[?](#)

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

Specifying Which Transporters Should be Used by the Replication Job

In NAKIVO Backup & Replication, Transporters perform all of the heavy lifting: read data from the source VM, compress the data, transfer the data over the network, perform data deduplication, and so on. In large and geographically distributed environments multiple Transporters can be deployed to distribute the data protection workload, optimize network traffic, and improve data transfer speeds. Thus, if more than one Transporter is deployed for NAKIVO Backup & Replication, it is important to determine which ones should be used to read data from a particular source VM and which ones should be used to write data to the target ESX(i) host and datastore.

By default, the product automatically determines which Transporter should be used to read data from the source VM. However, you can manually specify which Transporters should be used for the job. Here's how:

1. Click **Advanced options**.
2. Choose one of the following options:

Transporters

[?](#)

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved


- **Automatically determine source transporter:** The product will automatically determine which Transporters are the closest to source and target hosts.
- **Set source and target Transporters for all VMs:** Select this option to manually specify a single source and a single target Transporter that will be used for data transfer by the job.
- **Manually set transporters for source hosts:** Select this option to manually specify Transporters for all source and target hosts.

Completing the New Replication Job Wizard

Click **Finish** to complete the job creation.

Viewing Replication Job Information

Job information is aggregated in job dashboards. To open a dashboard for a particular job, click the job in the navigation pane:




Replication Job Widgets

Job details are available in job widgets. Refer to the sections below to learn about replication job widgets.

The “Replication Job Status” Widget

The “Replication Job Status” widget displays the current status of the job and the status of the last job run:



The “Source VMs” Widget

The “Source VMs” widget displays the list of VMs added to the replication job:

Source VMs		▼ Less details
2003forAmazonTest	Runs on demand	
Host:	192.168.77.213 > Ninja-Site 1 > 10.30.28.10	
Datastores:	esx10datastore	
Networks:	192.168.77.0	
testSY_win7	Runs on demand	
Host:	192.168.77.213 > Ninja-Site 1 > 10.30.28.10	
Datastores:	esx10datastore	
Networks:	192.168.77.0	

The following data is displayed for each VM:

- Power status of the source VMs
- Path to the host where each source VM is located
- Datastores where the source VM files are located
- Networks to which each source VM is connected

The “Replicas” Widget

The “Replicas” widget displays the list of replicas which are created (or will be created) by the job:

Replicas		▼ Less details
 amaz_Win7entSP1x32_SY_host77.9-re...	Not created yet	
Host:	10.30.30.40 > IT DC > esx243.zenlar.int	
Datastores:	esx243boot	
Networks:	10.30.30.0	

The following data is displayed:

- Names of replicas
- Host where the replica is located (or will be created)
- Datastore where the replica is located (or will be created)
- Networks to which the replica is (or will be) connected

The “Events” Widget

The “Events” widget displays all job-related events:

Events	
The job has finished The "New backup job" job has successfully finished.	Jul 12 at 12:07 PM
VM backup has started The "New backup job" job has started backing up the "2003forAmazonTest" VM.	Jul 12 at 12:07 PM
The job has started The "New backup job" job is starting its adhoc run. The job has run 1 time(s) previously.	Jul 12 at 12:07 PM
The job is running The "New backup job" job is currently running.	Jul 12 at 12:07 PM

The “Job Stats” Widget

The “Job Stats” widget displays statistical information about the job:


Job Stats	
 2 source VMs	
 4 disks	
 21 GB - size of source VMs	
 0.0% - data transfer compression on last run	
 0 minutes - duration of last run	
 0 minutes - average duration	

The following data is displayed:

- The number of source VMs replicated up by the job
- The number of disks replicated by the job
- Compression ratio achieved while transferring data over the network
- The duration of last run
- The average duration of job runs


The “Schedule and Options” Widget

The “Schedule and Options” widget displays the job schedule and replication options you have specified during the job creation:



The “Replica Datastores” Widget

The “Replica Datastores” widget displays all replica datastores to which the job replicates source VMs:




For each replica datastore, the following data is displayed:

- The amount of free and used space on the target datastore
- The size of the replicas

The “Average Bandwidth Utilization” Widget


The “Average Bandwidth Utilization” widget graphically displays the average data transfer speed for the last 10 job runs. The widget calculates and shows the average speed at which the data has been transferred. The information is based on the amount of time during which the job has completed and the amount of data transferred.



If one of the runs has failed, the widget displays data for that run as a red bar. If one of the runs has been stopped, the widget displays data for that run as a yellow bar.

The “Transferred Data” Widget


The “Transferred Data” widget graphically displays the amount of transferred data for the last 10 job runs.



If case one of the runs has failed, the widget displays data for that run as a red bar. If one of the runs has been stopped, the widget displays data for that run as a yellow bar.

The “Source Transporters” Widget


The source Transporters widget lists all Transporters that retrieve data from the source VMs and sends the data to the target Transporters:



Clicking a Transporter opens that Transporter in Configuration.

The “Target Transporters” Widget

The target Transporters widget lists all Transporters that receive data from the source Transporters and writes the data to the replica:



Clicking a Transporter opens that Transporter in Configuration.

Troubleshooting VM Replication

Replica VM Cannot be Powered On

NAKIVO Backup & Replication replicates almost all VM properties, including the MAC address. However, if this MAC address is out of standard VMware range (00:50:56:XX:XX:XX), you may not be able to power on the replica VM and will get the following message:

```
Failed to start the virtual machine. Module DevicePowerOn power on failed. Could not  
set up "macAddress" for ethernet0. Invalid MAC address specified. [MAC_address] is  
not an allowed static Ethernet address. It conflicts with VMware reserved MACs.
```

The workaround is described in the VMware knowledge base at:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2035707

Recovery

This section covers the following topics:

- [Recovering Files from VM Backups](#)
- [Recovering Microsoft Exchange Objects](#)
- [Recovering VMs from VM Backups](#)

Recovering Files from VM Backups


With NAKIVO Backup & Replication, you can recover files or folders directly from VM backups, without having to recover an entire VM first. File recovery is performed by the Transporter which is assigned to the backup repository from which you want to recover files.

NOTE: *File recovery is restricted to supported disk types and file systems on the backed up VMs. Refer to the [File Recovery Requirements](#) section for more information.*


Opening the File Recovery Wizard

To recover files or folders directly from VM backups, do either of the following:

- Open the File Recovery wizard from a job or a group dashboard by following the steps below:
 1. Select the backup job (or a group with such a job) that contains the backup from which you want to recover files.
 2. Click **Recover**, and then click **Recover Files**:




- Open the File Recovery wizard from the **Backup repositories** tab in **Configuration** by following the steps below:
 1. Click **Configuration** in the upper right corner of the product.
 2. Go to the **Backup repositories** tab and click a backup repository.
 3. In the backup repository title, click **Recover** and then click **Recover Files**.



The File Recovery wizard opens.

STEP 1: Choose a Backup and a Recovery Point

In the first step of the wizard, select a backup in the left pane and then select a recovery point in the right pane:




Click **Next** to go to the next step.

STEP 2: Choose Files to Recover

Searching for Files and Folders


To search for a file or a folder, type the full or a part of the item name in the **Search** field and press **Enter**:



NOTE: The search is performed starting from the point selected in the Navigation pane. For example, if you have selected **Hard drive 1 > Disk 1 > Program Files**, then the search will be performed only inside the **Program Files** folder.


Browsing Files and Folders

You can browse for files and folders of a VM backup by using the Navigation pane:




NOTE: If a VM backup contains Linux LVM volumes or Windows dynamic disks, the navigation pane will display these logical groups in addition to all hard drives available in the VM backup. If a hard drive does not contain any partitions and servers as a part of a Linux LVM volume or a Windows dynamic disk, such hard drive will appear as empty.

You can also quickly jump between folders by using the Speedbar:




Selecting Files and Folders to Recover

After you have located the item you'd like to recover, simply select the check box next to the item:



The number of items selected for recovery is displayed at the bottom of the wizard page. You can also:

- Click **Show** to view the list off all items selected for recovery
- Click **Clear selection** to clear the list of items selected for recovery



Forwarding Selected Files

After you have selected the items you want to recover, click **Forward selected**. In the view that appears, you can specify the **To** and **CC** fields, as well as email subject and body. The items you have selected for recovery will be added to the email as attachments.

The screenshot shows the NAKIVO Backup & Replication web interface. At the top, there's a navigation bar with 'Welcome, admin', 'Configuration', 'Support', 'Help', and 'Log out'. The main content area is titled 'File Recovery Wizard'. It has two steps: 'STEP 1: Choose a backup and a recovery point' and 'STEP 2: Choose files to recover'. The 'STEP 2' section is active and shows an email composition form. The form includes fields for 'To:', 'CC:', and 'Subject:'. The 'To:' and 'CC:' fields are filled with 'administrator@example.com'. The 'Subject:' field is filled with 'Recovered items - Mar 4 at 11:50 PM (UTC +02:00)'. Below these fields, there's an 'Attachments:' section showing two files: 'CatalogData-74ec826f-3e44-4751-a9a3-ccb56cb8728d-fe8cc6eb-e239-4965-aa3e-5823aff52375 (?)' and 'E00.chk (?)'. Below the attachments, there's a text area with the message 'Recovered by NAKIVO Backup & Replication'. At the bottom right, there are three buttons: 'Send', 'Cancel', and 'Close Wizard'. The footer of the interface says '© 2014 NAKIVO, Inc. All rights reserved'.

Downloading Selected Files

After you have selected the items you want to recover, click **Download**. All files and folders will now be sent to your browser for download as a single ZIP archive.

Troubleshooting File Recovery

Empty Disks in File Recovery Wizard

If you are receiving an empty disk when trying to perform file recovery, this may be due to:

- In case the Transporter that is assigned to the backup repository from which you are trying to recover files is installed on RHEL, try to disable SELINUX by following the steps below:
 - n. Open the config by running the following command:


```
vi /etc/selinux/config
```
 - o. Change the "SELINUX=" value from "enforcing" to "disabled".
 - p. save the changes by typing the **Escape** key, then typing:


```
:wq
```
 - q. Restart the machine on which the Transporter is installed.
- The disk is a part of RAID; the current version of the product does not support file recovery from RAID disks.
- The disk is a part of LVM; the current version of the product does not support file recovery from LVM disks.

- The disk does not have a partition table; the current version of the product cannot work with disks which do not have a partition table.
- The disk has an unsupported partition table; the current version of the product cannot work with proprietary disk partition tables. Only the MBR and GPT partition tables are supported.

If the above is not the case, the error might be caused by a high load on the Transporter that is assigned to the backup repository. This can occur if numerous VM disks are currently being processed by this Transporter.

To fix this problem, please make sure the Transporter assigned to the repository is not overloaded. You can simply stop the jobs that use this Transporter and try to recover the files again. As an alternative, you can decrease the maximum Transporter load by following the steps below:

1. Open **Configuration > Transporters**.
2. **Edit** the Transporter that is overloaded.
3. Decrease the Maximum Load.
4. Save the Transporter.
5. Wait until the actual Transporter load will be decreased. This can take some time.
6. Try recovering files again.

Unsupported File System for File Recovery

It might be that you are trying to recover files from a Linux VM using a Transporter installed on Windows. Windows OS does not support Linux file systems, so the Transporter installed on Windows cannot recover files from Linux. To resolve this issue, follow the steps below:

1. Install a Transporter on a Linux OS.
2. Assign the Transporter to the backup repository from which you try to recover files.
3. Try recovering files again.

If your Transporter is installed on Linux and a partition was reported as unsupported, this may be due to:

- The Transporter does not have drivers for this partition file system. Install a proper file system driver on the Transporter machine and try recovering files again.
- The partition is a raw partition and does not have a file system. The current version of the product does not support raw partitions.


Recovering Microsoft Exchange Objects

NAKIVO Backup & Replication enables browsing, searching, and recovering Microsoft Exchange emails directly from compressed and deduplicated VMware VM backups, without the need to restore the entire VM first. This feature is agentless, works right out of the box, and does not require creating a special lab or running a special backup type.


Opening the Exchange Object Recovery Wizard

To start recovering Microsoft Exchange objects from VM backups, do either of the following:

- Open the Exchange Object Recovery Wizard from a job or a group dashboard by following the steps below:
 - a. Select the backup job or a group that contains the backup of a VM running Microsoft Exchange.
 - b. Click **Recover**, and then click **Recover Exchange Objects**




- Open the Exchange Object Recovery Wizard from the **Backup repositories** tab in **Configuration**:
 - a. Click **Configuration** in the upper right corner of the product.
 - b. Go to the **Backup repositories** tab and click a backup repository.
 - c. Click a VM backup, and then choose **Recover Exchange Objects** from the **Recover** menu:



The New Recovery Job Wizard opens.

STEP 1: Choose a Backup and a Recovery Point

In the first step of the wizard, select a backup of a VM with Microsoft Exchange in the left pane and then select a recovery point in the right pane:



Click **Next** to go to the next step.

Manually locating Exchange Database Files

By default, NAKIVO Backup & Replication automatically searches the selected recovery point for Microsoft Exchange databases (files with “.edb” extension) from which application objects can be recovered. This process can take a few minutes. If you want to manually specify the location of the database file, deselect the **Automatically locate application databases** option.


STEP 2: Choose Application Items to Recover

Searching for Microsoft Exchange Objects

NAKIVO Backup & Replication provides the ability to search for emails by email subject or body. The search functionality, however, has the following limitations:

- The product can search only for emails by email subject or email body
- If text formatting (such as “bold text”) is applied to a keyword that is searched for, the search may not find the keyword due to formatting conversion issues.
- The product does not create or maintain an index of the Exchange database contents. The search is performed on the fly and can take a long time to complete.
- To speed up the search, perform the search within a particular folder, rather than in a mailbox.

To search for an email by its subject or body, type a word in the **Search** field and press **Enter**:




NOTE: The search is performed starting from the point selected in the left (Navigation) pane. For example, if you have selected Mailbox Database > John Smith, then the search will be performed only inside the John Smith mailbox.

Browsing Exchange Objects


NAKIVO Backup & Replication scans the selected recovery point for Microsoft Exchange databases (files with “.edb” extension), and displays the list of found databases in the left (Navigation) pane.

NOTE: Not all of the found database files contain Microsoft Exchange objects that can be recovered by the product.

To browse Microsoft Exchange objects, simply expand the appropriate database in the left pane:



You can also jump in the tree by using the Speedbar:




Viewing Microsoft Exchange Objects

To view a Microsoft Exchange object – such as an email – simply click the object. Object contents will be displayed.


Selecting Microsoft Exchange Objects to Recover

In the Contents pane, select check boxes next to files and folders you want to recover:



The number of items selected for recovery is displayed at the bottom of the wizard page. You can also:

- Click **Show** to view the list off all items selected for recovery
- Click **Clear selection** to clear the list of items selected for recovery



Forwarding Selected Objects

NAKIVO Backup & Replication provides the ability to send the objects you have selected for recovery over email right from the product interface.

After you have selected the items you want to recover, click **Forward selected**. In the view that appears, you can specify the **To** and **CC** fields, as well as email subject and body. The items you have selected for recovery will be added to the email as attachments.

Downloading Selected Objects

After you have selected the objects you want to recover, click **Download**. All objects will be sent to your browser for download as a single ZIP archive.

Troubleshooting Exchange Objects Recovery

Exchange Database Cannot be Found

If the Exchange Object Recovery wizard cannot find the Exchange database, do the following:

- Make sure that your Exchange version is supported. Refer to the [Feature Requirements](#) section for more information.
- Locate the Exchange database manually as described in the [Manually locating Exchange Database Files](#) section.


Recovering VMs from VM Backups

With NAKIVO Backup & Replication, you can recover full VMs from backups. When you run VM recovery, a new VM is created; the source VM is not reverted to a previous state or replaced with the new VM.


Opening the New Recovery Job Wizard

To recover full VMs from backups, do either of the following:

- Open the New Recovery Job Wizard from a job or a group dashboard by following the steps below:
 1. Select the backup job (or a group with such a job) that contains the backup from which you want to recover a VM.
 2. Click **Recover**, and then click **Recover VMs**:




- Open the New Recovery Job Wizard from the **Backup repositories** tab in **Configuration** by following the steps below:
 1. Click **Configuration** in the upper right corner of the product.
 2. Go to the **Backup repositories** tab and click a backup repository.
 3. In the backup repository title, click **Recover** and then click **Recover VMs**.



The New Recovery Job Wizard opens.

STEP 1: Choose VMs to Recover

In the first step of the wizard, select one or more VM backups in the left pane and then select a recovery point for each backup in the right pane:




Click **Next** to go to the next step.

STEP 2: Choose Recovery Location

Setting the Same Host, Datastore, and Network for All Recovered VMs

To recover all VMs to the same container and datastore as well as connect all recovered VMs to the same networks, follow the steps below:

1. Choose a cluster, host, or resource pool from the **Target** dropdown menu.
2. Choose a datastore from the **Datastore** dropdown menu.
3. Choose a network from the **Network** dropdown menu.



Setting Different Options for Recovered VMs

To specify different options for recovered VMs, follow the steps below:

1. Click **Advanced setup**.
2. Choose a target host, target datastore, and target network for each VM:

NAKIVO Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Recovery Job Wizard

STEP 1: Choose VMs to recover

STEP 2: Choose recovery location

Recovery from backup

Target: 10.30.28.23 | Datastore: esx21datastore | Network: 10.30.30.0

VM:	Host:	Path:
exchange01.corp.int	10.30.28.23	10.30.29.12 > IT
1 disk:		esx21datastore (432.2 GB free)
1 NIC:		10.30.30.0

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

3. Click **Next** to go to the next step.

STEP 3: Specify Recovery Options

Setting general options

Specify the general options as follows:

NAKIVO Backup & Replication

Welcome, admin | Configuration | Support | Help | Log out

New Recovery Job Wizard

STEP 1: Choose VMs to recover

STEP 2: Choose recovery location

STEP 3: Specify recovery options

General options

Job name: New recovery job

☒ Run job immediately after creation

☒ Power on recovered VMs

☐ Generate new MAC addresses for recovered VMs

Recovery from backup

☒ Synthetic recovery

☐ Full production recovery

☐ Enable network acceleration

☐ Encrypt data during transfer

☐ Create only thin disks on recovered VMs

Advanced options...

Next Close Wizard


© 2014 NAKIVO, Inc. All rights reserved

- **Job name:** specify a name for the recovery job.

- **Run job immediately after creation:** if this option is selected, the job will be started automatically after you complete the wizard.
- **Power on recovered VMs:** If this option is selected, the recovered VMs will be powered on.
- **Generate new MAC address for recovered VMs:** If this option is deselected, the recovered VM will have the same MAC address as the source VM. If this option is selected, a new MAC address will be generated for the recovered VM.

Specify the Recovery Options

Specify the recovery options as follows:



- **Synthetic recovery:** If this option is selected, the VMs will be recovered with the environmental dependencies (such as CPU Affinity) removed. Select this option when recovering VMs to a new location.
- **Full production recovery:** If this option is selected, the recovered VMs will be identical to the backed up VMs. Select this option to recover VMs to their original location.
- **Enable network acceleration:** If this option is selected, NAKIVO Backup & Replication will use compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to recover VMs over WAN or slow LAN links.
- **Encrypt data during transfer:** If this option is selected, VM data will be protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on machines running Transporters. Select this option if recover over WAN without a VPN connection.
- **Create only thin disks for recovered VMs:** If this option is selected, only thin disks will be created on recovered VMs, regardless of the disk type of the original source VM.

Specifying VM Replica Names

NAKIVO Backup & Replication provides the ability to change recovered VM names so that the recovered VMs can be easily distinguished from the source VMs. By default, the text “- recovered” is appended to the end of the recovered VM name.

To change VM Replica names, do the following:

1. Click **Advanced options**.
2. Choose between the following options:

STEP 3: Specify recovery options

General options

Job name:

☒ Run job immediately after creation

☒ Power on recovered VMs

☐ Generate new MAC addresses for recovered VMs

Recovery from backup

☒ Synthetic recovery ?

☐ Full production recovery ?

☐ Enable network acceleration ?

☐ Encrypt data during transfer ?

☐ Create only thin disks on recovered VMs ?

Recovered VM names

☒ Append VM name

☐ Leave VM names as is

☐ Enter custom VM names

Post job actions

☒ Send job reports on each job completion to: ?

Transporters

?

Next Close Wizard

© 2014 NAKIVO, Inc. All rights reserved

- **Append X before/after VM name** – source VM names will be used for recovered VM names and the specified text will be added before or after the recovered VM name.
- **Leave VM names as is** – recovered VM names will be identical to the source VM names.
- **Enter custom replica names** – allows you to enter custom names for recovered VM.

Setting up Email Notifications for the Job

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications described in the [Configuring Email Notifications](#) section and provides the ability to configure notifications on the per-job level.

NOTE: To enable this option, configure your Email settings as described in the [Configuring Email Settings](#) section.

To send email notifications, do the following:

1. Click **Advanced options**.
2. Select the **Send job reports on each job completion to** option and specify one or more email addresses in the text field. The semi-colon character should be used to separate multiple email addresses.

STEP 3: Specify recovery options

General options

Job name:

☒ Run job immediately after creation

☒ Power on recovered VMs

☐ Generate new MAC addresses for recovered VMs

Recovery from backup

☒ Synthetic recovery ?

☐ Full production recovery ?

☐ Enable network acceleration ?

☐ Encrypt data during transfer ?

☐ Create only thin disks on recovered VMs ?

Recovered VM names

☒ Append VM name

☐ Leave VM names as is

☐ Enter custom VM names

Post job actions


☒ Send job reports on each job completion to: ?

Specifying Which Transporters Should be Used by the Replication Job

In NAKIVO Backup & Replication, Transporters perform all of the heavy lifting: read data from the source VM, compress the data, transfer the data over the network, perform data deduplication, and so on. In large and geographically distributed environments multiple Transporters can be deployed to distribute the data protection workload, optimize network traffic, and improve data transfer speeds. Thus, if more than one Transporter is deployed for NAKIVO Backup & Replication, it is important to determine which ones should be used to read data from a particular source VM and which ones should be used to write data to the target ESX(i) host and datastore.

By default, the product automatically determines which Transporter should be used to read data from the source VM. However, you can manually specify which Transporters should be used for the job. Here's how:

1. Click **Advanced options**.
2. Choose one of the following options:



- **Automatically determine source transporter:** The product will automatically determine which Transporters are the closest to source and target hosts.
- **Set source and target Transporters for all VMs:** Select this option to manually specify a single source and a single target Transporter that will be used for data transfer by the job.
- **Manually set transporters for source hosts:** Select this option to manually specify Transporters for all source and target hosts.

Completing the New Recovery Job wizard

Click **Finish** to complete the job creation. If the **Run job immediately after creation** option has been selected, the recovery job will run immediately. You can also run the job by clicking **Run Job** in the recovery job dashboard.

Job Management

This section covers the following topics:

- [Job Alarms and Notifications](#)
- [Running Jobs on Demand](#)
- [Stopping Jobs](#)
- [Editing Jobs](#)
- [Disabling and Enabling Jobs](#)
- [Grouping Jobs](#)

Grouping Jobs


Groups are folders which allow you to:

- Logically arrange jobs (to represent organizations, locations, services, etc.).
- Perform bulk actions with all or selected jobs in a group.

Creating Groups

To create a group, follow the steps below:

1. Click **Create** and then click **Group** on any dashboard:



2. Type in the group name in the dialog box that appears and click **OK**.

Managing Groups

The following actions are available to manage groups:

- To add a job to a group, simply drag the job into the group.
- To remove a job from the group, drag the job outside the group.

- To delete a group, right-click the group and choose **Delete** from the shortcut menu that appears. Confirm the group deletion when prompted to do so. Note that when deleting a group, its jobs are not deleted and are moved to the parent group (or Overview).
- To rename a group, double-click the group and enter a new name.
- To enable or disable all jobs inside a group, click the **Enable/Disable** switch.
- To run jobs available in a group, click **Run/Stop** and then click **Run Jobs**. In the dialog box that appears, select the jobs you want to run and click **Run Jobs**.
- To stop running jobs available in a group, click **Run/Stop** and then click **Stop Jobs**. In the dialog box that appears, select the jobs you want to stop and click **Stop Jobs**.
- Deleting Jobs


Job Alarms and Notifications

NAKIVO Backup & Replication displays:

- **Alarms:** backup failures
- **Notifications:** infrastructure changes and job errors that do not lead to backup failure

Viewing Job Alarms and Notifications

Alarms and notifications are displayed below the job name. If there are more than five alarms in a job, you can jump between pages by clicking the navigation controls:




Dismissing Job Alarms

To dismiss all alarms and notifications in a job, click **Dismiss All**. To dismiss an individual alarm or notification, hover the mouse pointer over the alarm or notification and click **Dismiss**.

Running Jobs on Demand

To start a job, follow the steps below:

1. Go to the job's dashboard and click **Run Job**:




2. In the message that appears, click **Run Job**.

Stopping Jobs

To stop a job that is currently running, follow the steps below:

1. Go to the job's dashboard and click **Stop Job**:




2. Click **Stop Job** in the message that appears.

Editing Jobs

To edit a job, follow the steps below:


1. On the job dashboard, click **Manage** and then click **Edit Job**:



2. In the Edit wizard, click on a step where you want to make changes.
3. Make the required changes and then click **Save and Close**.

Disabling and Enabling Jobs

NAKIVO Backup & Replication allows the ability to disable jobs. A disabled job does not run on a schedule, nor can it be run on demand. You can disable or enable jobs at any time using the **Enable/Disable** switch:



Grouping Jobs


Groups are folders which allow you to:

- Logically arrange jobs (to represent organizations, locations, services, etc.).
- Perform bulk actions with all or selected jobs in a group.

Creating Groups

To create a group, follow the steps below:

4. Click **Create** and then click **Group** on any dashboard:



5. Type in the group name in the dialog box that appears and click **OK**.

Managing Groups


The following actions are available to manage groups:

- To add a job to a group, simply drag the job into the group.
- To remove a job from the group, drag the job outside the group.
- To delete a group, right-click the group and choose **Delete** from the shortcut menu that appears. Confirm the group deletion when prompted to do so. Note that when deleting a group, its jobs are not deleted and are moved to the parent group (or Overview).
- To rename a group, double-click the group and enter a new name.
- To enable or disable all jobs inside a group, click the **Enable/Disable** switch.
- To run jobs available in a group, click **Run/Stop** and then click **Run Jobs**. In the dialog box that appears, select the jobs you want to run and click **Run Jobs**.
- To stop running jobs available in a group, click **Run/Stop** and then click **Stop Jobs**. In the dialog box that appears, select the jobs you want to stop and click **Stop Jobs**.

Deleting Jobs

To delete a job follow the steps below:

1. On the job dashboard, click **Manage** and then click **Delete Job**:



2. Click **Delete** Job in the message that appears.

Cloud Integration

This section covers the following topics:

- [Amazon EC2 Integration](#)
- [vCloud Director Integration](#)

Amazon EC2 Integration

NAKIVO Backup & Replication allows an automated integration with Amazon EC2, a cost-effective public cloud. Backing up to the cloud is as simple as backing up to a regular backup repository.

Setting up Integration with Amazon

To start backing up to the Amazon EC2, you must first register with Amazon Web Services and obtain the AWS Security Credentials.

To integrate with Amazon cloud, follow the steps below:

1. In the right upper corner of NAKIVO Backup & Replication, click **Configuration**.
2. Open the **Cloud Integration** tab.
3. Click the **New Cloud Integration** button.
4. Fill in the **AWS Access Key ID field**. Enter the AWS Access Key ID that is issued when you create your AWS account.
5. Fill in the **Secret Access Key field**. Enter the Secret Access Key that is issued when you create your AWS account.
6. Fill in the **Region** field. Choose the Region that is closest to your production environment to receive higher bandwidth. The region cannot be changed later.
7. Fill in the **Storage (GB)** field. This is the initial amount of backup storage that will be allocated in the cloud. This value can be changed later.
8. Fill in the **Storage chunk** field. Amazon provides Amazon Elastic Block Store (EBS) block level storage volumes for use with Amazon EC2 instances. When creating a backup repository, the product creates several EBS volumes and represents them as a single storage. The **Storage chunk** field specifies the size of each EBS volume. The product can create up to 50 EBS volumes per integration.
9. Fill in the **Inbound firewall rules** field. Specify one or more CIDR ranges to be allowed to connect to your Amazon instance. You can leave this field by default if you are not sure.
10. To change the default port that is used by the Director to communicate with the Onboard Transporter, enter a new value in the **Transporter port** field.
11. Specify a range of port numbers (from 1 to 65535) that will be used to transfer data in the **Data transfer ports** field. The range you specify should contain at least 100 ports.
12. Select the following additional options:
 - **Automatically resize storage**: Selecting this option will enable the automatic storage resize. The product will automatically:
 - Increase the storage size by the value specified in the **Storage chunk** field if there's less than 10% of free space left in the backup repository.
 - Decrease the storage size by the value specified in the **Storage chunk** field if the free space is more than 200% of the value specified in the **Storage chunk** field and there is more than 15% of free space in the backup repository.
 - **Compress backups in cloud**: Selecting this option will enable storing all backups created by this integration in a compressed form. This option cannot be changed later.
 - **De-duplicate backups in cloud**: Selecting this option will enable storing all backups created by this integration in a de-duplicated form. This option cannot be changed later.

- **Encrypt backups in cloud:** Selecting this option will enable storing all backups created by this integration in an encrypted form. All data uploaded to the Amazon cloud will be automatically encrypted with the AES 256-bit encryption. This option cannot be changed later.

13. Click the **Add** button to start the automated integration. The process takes approximately 10 minutes.

Managing Amazon Integration

Downloading Amazon Integration Keys

We recommend downloading the integration keys in a zip archive. They will be essential to import this integration later.

Editing Amazon Integration

1. In the right upper corner of NAKIVO Backup & Replication, click the **Configuration** link.
2. Open the **Cloud Integration** tab.
3. Hover your mouse over a specific Amazon Integration and click the **Edit** link.
4. You can make changes to the following fields:
 - AWS Access Key ID
 - Secret Access Key
 - Storage (GB)
 - Inbound connections

The **Automatically resize storage** option can also be changed. However, the Region, Compress backups in cloud, De-duplicate backups in cloud, and Encrypt backups in cloud options cannot be changed.

Click the **Save** button to finish the operation. NAKIVO Backup & Replication will verify the changes and apply them.

Removing Amazon Integration

1. In the right upper corner of NAKIVO Backup & Replication, click the **Configuration** link.
2. Open the **Cloud Integration** tab.
3. Hover your mouse over a specific Amazon Integration and click the **Remove** link. You will receive a warning message with two available options:
 - **Remove Integration and Keep Cloud Objects:** If this option is selected, you will be able to import this integration later. To perform a successful import, you will need to have keys for this integration.
 - **Remove Integration and Delete Cloud Objects:** If this option is selected, all backups and cloud objects created by this integration will be removed.

Importing Existing Amazon Integration

1. In the right upper corner of NAKIVO Backup & Replication, click the **Configuration** link.
2. Open the **Cloud Integration** tab.
3. Click the **Import Integration** button.
4. Fill in the **AWS Access Key ID** field.
5. Fill in the **Secret Access Key** field.
6. Click the **Import** button. NAKIVO Backup & Replication will search for the existing integrations using provided security credentials.
7. Once your integrations are found, click the **Download keys** link to upload the archive with keys.

8. Select the check boxes next to specific integrations you need to import.
9. Click the **Import selected** button to finish the operation.

vCloud Director Integration

NAKIVO Backup & Replication can now back up and replicate vCloud Director vApp VMs, recover files from vApp VM backups, as well as recover full VMs back into vApps.

Backing Up and Replicating vCloud Director VMs

To back up vApp VMs, follow the steps below:

1. Discover the vCenter that is used by your vCloud Director as described in the [Adding vCenters and ESX\(i\) hosts](#) section.
2. In the vSphere inventory, vApps are represented as resource pools. In NAKIVO Backup & Replication, you can locate the resource pools that represent vApps and then select either entire (vApp) resource pools or individual (vApp) VMs as described in the [Adding VMs and Containers to the Backup Job](#) and [Adding VMs and Containers to the Replication Job](#) sections.

During the job run, the product will backup/replicate all VM attributes related to vCloud Director so that the VM can be recovered back to vCloud Director vApp.

Recovering vCloud Director VMs back to vApps

To recover backed up VMs back to vApps, follow the steps below:

1. Remove the source VM from the (vApp) resource pool in vSphere so there is no conflict when you recover the VM from the backup.
2. Create a recovery job as described in the [Empty Disks in File Recovery Wizard](#)
3. If you are receiving an empty disk when trying to perform file recovery, this may be due to:
4. In case the Transporter that is assigned to the backup repository from which you are trying to recover files is installed on RHEL, try to disable SELINUX by following the steps below:
5. Open the config by running the following command:
6. `vi /etc/selinux/config`
7. Change the "SELINUX=" value from "enforcing" to "disabled".

save the changes by typing the Escape key, then typing:

- `:wq`
- Restart the machine on which the Transporter is installed.
- The disk is a part of RAID; the current version of the product does not support file recovery from RAID disks.

The disk is a part of LVM; the current version of the product does not support file recovery from LVM disks.

The disk does not have a partition table; the current version of the product cannot work with disks which do not have a partition table.

The disk has an unsupported partition table; the current version of the product cannot work with proprietary disk partition tables. Only the MBR and GPT partition tables are supported.

- If the above is not the case, the error might be caused by a high load on the Transporter that is assigned to the backup repository. This can occur if numerous VM disks are currently being processed by this Transporter.
 - d. To fix this problem, please make sure the Transporter assigned to the repository is not overloaded. You can simply stop the jobs that use this Transporter and try to recover the files again. As an alternative, you can decrease the maximum Transporter load by following the steps below:
 - e. Open Configuration > Transporters.
 - f. Edit the **Transporter that is** overloaded.

Decrease the Maximum Load.

- Save the Transporter.
- Wait until the actual Transporter **load** will be decreased. This can take some time.
 - g. Try **recovering** files again.
 - h. Unsupported File System for File Recovery
 - i. It might be **that** you are trying to recover files from a Linux VM using a Transporter installed on Windows. Windows OS does not support Linux file systems, so the Transporter installed on Windows cannot recover files from Linux. To resolve this issue, follow the steps below:
 - j. Install a Transporter on a Linux OS.
 - k. Assign the **Transporter** to the backup repository **from which** you try to recover files.

Try recovering files again.

If your Transporter is installed on Linux and a partition was reported as unsupported, this may be due to:

The Transporter does not have drivers for this partition file system. Install a proper file system driver on the Transporter machine and try recovering files again.


The partition is a raw partition and does not have a file system. The current version of the product does not support raw partitions.

Recovering Microsoft Exchange Objects

1. NAKIVO Backup & Replication enables browsing, searching, and recovering Microsoft Exchange emails directly from compressed and deduplicated VMware VM backups, without the need to restore the entire VM first. This feature is agentless, works right out of the box, and does not require creating a special lab or running a special backup type.

Opening the Exchange Object Recovery Wizard


2. To start recovering Microsoft Exchange objects from VM backups, do either of the following:
3. Open the Exchange Object Recovery Wizard from a job or a group dashboard by following the steps below:
4. Select the backup job or a group that contains the backup of a VM running Microsoft Exchange.
5. Click **Recover**, and then click **Recover Exchange Objects**



- 6.
7. **Open** the Exchange Object Recovery Wizard from the **Backup repositories** tab in **Configuration**: Click **Configuration** in the upper right corner of the product.

Go to the Backup repositories tab and click a backup repository.


Click a VM backup, and then choose **Recover Exchange Objects** from the **Recover** menu:



The New Recovery Job Wizard opens.

STEP 1: Choose a Backup and a Recovery Point

In the first step of the wizard, select a backup of a VM with Microsoft Exchange in the left pane and then select a recovery point in the right pane:



- Click Next to go to the next step.
- Manually locating Exchange Database Files
- By default, NAKIVO Backup & Replication automatically searches the selected recovery point for Microsoft Exchange databases (files with “.edb” extension) from which application objects can be recovered. This process can take a few minutes. If you want to manually specify the location of the database file, deselect the **Automatically locate application databases** option.
- STEP 2: Choose Application Items to Recover
- Searching for Microsoft Exchange Objects
- NAKIVO Backup & Replication provides the ability to search for emails by email subject or body. The search functionality, however, has the following limitations:


The product can search only for emails by email subject or email body

If text formatting (such as “bold **text**”) is applied to a keyword that is searched for, the search may not find the keyword due to formatting conversion issues.

The product does not create or maintain an index of the Exchange database contents. The search is performed on the fly and can take a long time to complete.

To speed up the search, perform the search within a particular folder, rather than in a mailbox.

To search for an email by its subject or body, type a word in the Search field and press Enter:




NOTE: The search is performed starting from the point selected in the left (Navigation) pane. For example, if you have selected Mailbox Database > John Smith, then the search will be performed only inside the John Smith mailbox.

Browsing Exchange Objects


NAKIVO Backup & Replication scans the selected recovery point for Microsoft Exchange databases (files with ".edb" extension), and displays the list of found databases in the left (Navigation) pane.

NOTE: Not all of the found database files contain Microsoft Exchange objects that can be recovered by the product.

To browse Microsoft Exchange objects, simply expand the appropriate database in the left pane:



You can also jump in the tree by using the Speedbar:




Viewing Microsoft Exchange Objects

To view a Microsoft Exchange object – such as an email – simply click the object. Object contents will be displayed.

Selecting Microsoft Exchange Objects to Recover


In the Contents pane, select check boxes next to files and folders you want to recover:



The number of items selected for recovery is displayed at the bottom of the wizard page. You can also:

Click **Show** to view the list off all items selected for recovery

Click **Clear selection** to clear the list of items selected for recovery



Forwarding Selected Objects

NAKIVO Backup & Replication provides the ability to send the objects you have selected for recovery over email right from the product interface.

- After **you** have selected the items you want to recover, click Forward **selected**. In the view that appears, you can specify the **To** and **CC** fields, as well as email subject and body. The items you have selected for recovery will be added to the email as attachments.
- Downloading **Selected Objects**
- After you have selected the objects you want to recover, click **Download**. All objects will be sent to your browser for download as a single ZIP archive.

Troubleshooting Exchange Objects Recovery

Exchange Database Cannot be Found

If the Exchange Object Recovery wizard cannot find the Exchange database, do the following:

Make sure that your Exchange version is supported. Refer to the [Feature Requirements](#) section for more information.

Locate the Exchange database manually as described in the Manually locating Exchange **Database** Files section.

8. Recovering VMs from VM Backups section and make sure you select the **Full production recovery** option as described in the [Specify the Recovery Options](#) section.

Reporting


This section covers the following topics:

- [Setting up Email Reporting](#)
- [Generating Reports on Demand](#)

Setting up Email Reporting

NAKIVO Backup & Replication can send notifications and reports over email. To configure the automatic reports, follow the steps below:


1. Click **Configuration** in the upper right corner of the product and select to the **General** tab.
2. Make sure you have configured email settings as described in the [Configuring Email Settings](#) section.
3. Click **Automatic reports**.
4. In the title of the **Automatic reports** box, click **Edit**:



5. Select or deselect the following email notification options:
 - **Send job reports on each job** – if this option is selected, NAKIVO Backup & Replication will send an HTML report after completion of every job (regardless of the job success or failure) to email addresses specified in the text field. Use a semi-colon to separate multiple email addresses.
NOTE: Changing this setting will affect only new jobs; existing jobs should be updated manually.
 - **Send notifications on errors to** – if this option is selected, NAKIVO Backup & Replication will send an email notification on job, repository, infrastructure, connection, and other failures to email addresses specified in the text field. Use a semi-colon to separate multiple email addresses.
 - **Send notifications on warnings to** – if this option is selected, NAKIVO Backup & Replication will send an email notification on non-critical events, such as infrastructure change, to email addresses specified in the text field. Use a semi-colon to separate multiple email.
 - **Send Overview report on schedule to job** – if this option is selected, NAKIVO Backup & Replication will generate the Overview report (which includes information about all jobs in the product) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semi-colon to separate multiple email addresses.
6. Click **Apply**.

Generating Reports on Demand

To generate a report, simply click **Create** and then click **Report** on any dashboard:



Command Line Interface

This section covers the following topics:

- [Using Command Line Interface](#)
- [Available Commands and Exit Codes](#)

Using Command Line Interface

Using Command Line Interface Locally

To use the product's command line interface (CLI) on the machine where NAKIVO Backup & Replication Director is installed, follow the steps below:

1. Run the CLI executable:
 - If NAKIVO Backup & Replication is installed on a Windows OS, run the **cli.bat** file located in the **bin** folder inside the product installation folder ("C:\Program Files\NAKIVO Backup & Replication" by default).
 - If NAKIVO Backup & Replication is installed on a Linux OS, run the **cli.sh** file located in the **director/bin** folder inside the product installation folder ("/opt/nakivo/" by default).
2. Run commands described in the [Available Commands](#) section.

Using Command Line Interface Remotely

To use the product's command line interface (CLI) from a remote machine, follow the steps below:

1. Copy the CLI executable and jar files to the machine from where you plan to use the CLI:
 - If NAKIVO Backup & Replication is installed on a Windows OS, copy the **cli.bat** and **cli.jar** files located in the **bin** folder inside the product installation folder ("C:\Program Files\NAKIVO Backup & Replication" by default).
 - If NAKIVO Backup & Replication is installed on a Linux OS, copy the **cli.sh** and **cli.jar** files located in the **director/bin** folder inside the product installation folder ("/opt/nakivo/" by default).
3. On the machine from where you plan to use the CLI, configure the PATH system variable as described at <http://java.com/en/download/help/path.xml>
4. Run commands using the following format:

```
<command> <host> <port> <username> <password>
```

Example: To get a list of jobs of the product which is installed on the machine with the 192.168.10.10 IP address, uses the 4443 port number for the **Director Web HTTPS port**, and has "admin" as login and password for the product's web UI, run the following command:

```
--job-list --host 192.168.10.10 --4443 --username admin --password admin
```

Available Commands and Exit Codes

Available Commands

Command	Argument	Description
-? --?	none	Print help in the console
-f --save-to	Local path	Set the default path for saving reports and support bundles
-h --help	none	Print help in the console
-i --interactive	none	Enter interactive mode
-jd --job-disable	Job ID	Disable the job specified by the argument
-je --job-enable	Job ID	Enable the job specified by the argument
-ji --job-info	Job ID	Print job info
jl --job-list	none	List all jobs
-jp --job-report	Job ID	Create job/overview report
-jr --job-start	Job ID	Run the job specified in the argument
-js --job-stop	Job ID	Stop the job specified in the argument
-p --port	Web UI port number	Use the port specified in the argument
-pwd --password	Web UI password	Password used for Web UI login
-ra	Repository ID	Attach backup repository

--repository-attach		
-rc --repository-consolidate	Repository ID	Start backup repository consolidation
-rcs --repository-consolidate-stop	Repository ID	Stop backup repository consolidation
-rd --repository-detach	Repository ID	Detach backup repository
-ri --repository-info	Repository ID	Print backup repository information
-rl --repository-list	none	List all backup repositories
-ru --repository-update	Repository ID	Run update of all backup repositories
-t --host	IP of the host	Host where the product is installed
-u --username	Web UI username	Use the specified username

Exit Codes

NAKIVO Backup & Replication CLI provides the following exit codes:

- **0:** Normal
- **1:** Unknown command
- **2:** Cannot login
- **3:** Command failed
- **4:** Local failure
- **5:** No arguments

Troubleshooting

This section covers the following topics:

- [Generic Disk-Level Job Failure](#)
- [Generic VM-Level Job Failure](#)
- [VM Quiescing Failure](#)
- [Empty Disks in File Recovery Wizard](#)
- [Unsupported File System for File Recovery](#)
- [Enabling CBT for a VM](#)
- [Enabling SSH in ESXi Hosts](#)

Generic Disk-Level Job Failure

Refer to the sections below to learn about possible causes of the following errors:

- “Backup/replication/recovery of the VM has failed”
- “An error has occurred while backing up/replicating/recovering one or more disks of the VM.”

Network or environment configuration problems

Invalid IP/DNS Configuration

There can be problems with resolving source or target host DNS names on the machine where the Transporter is installed. Also, the source or target host can have multiple IP addresses in different networks, and some of them might not be reachable from the Transporter. Please make sure the source and/or target host can be pinged from the appropriate Transporter machine. If ping fails, please try to edit the HOSTS file on the Transporter machine and specify the proper DNS name and IP address.

Closed Ports

The product requires a number of TCP ports to be open in order to operate successfully. These ports include:

- Director Web HTTPS port (4443 by default): used to access the Web UI of the product. It must be open on the machine where the product is installed.
- Transporter Port (9446 by default): used for communication with Transporters. It must be open on every machine where the Transporter is installed (including the Director one).
- Data transfer ports (9447-10000 by default): ports from this range are used for the VM data transfer. This range must be open on every machine where the Transporter is installed (including the Director one).
- VMware Web Services Port (443 by default): used to access VMware infrastructure. It must be open on source and target hosts used for VM backup, replication, or recovery.
- VMware Data Transfer Port (902 by default): used to access VMware infrastructure. It must be open on source and target hosts used for VM backup, replication, or recovery.

Please make sure all the described ports are not blocked by firewall, as this may be the cause of the above error.

Unexpected Server Outage During the Job Run

Source or Target host outage

During the job execution, the source or target host has been disconnected or powered off.

Please make sure the source host and/or the target host is powered on and accessible via network.

Source or Target Transporter outage

During the job execution, the source or target Transporter has been disconnected or powered off.

Please make sure the source host and/or the target host is powered on and accessible via network.

Generic VM-Level Job Failure

Refer to the sections below to learn about possible causes of the following errors:

- “Backup/Replication/Recovery of the VM has failed”
- “An internal error has occurred while backing up/replicating/recovering the VM.”

Network or Environment Configuration Problems

Closed Ports

The product requires a number of TCP ports to be open in order to operate successfully. These ports include:

- Director Web HTTPS port (4443 by default): used to access the Web UI of the product. It must be open on the machine where the product is installed.
- Transporter Port (9446 by default): used for communication with Transporters. It must be open on every machine where the Transporter is installed (including the Director one).
- Data transfer ports (9447-10000 by default): ports from this range are used for the VM data transfer. This range must be open on every machine where the Transporter is installed (including the Director one).
- VMware Web Services Port (443 by default): used to access VMware infrastructure. It must be open on source and target hosts used for VM backup, replication, or recovery.
- VMware Data Transfer Port (902 by default): used to access VMware infrastructure. It must be open on source and target hosts used for VM backup, replication, or recovery.

Please make sure all the described ports are not blocked by firewall, as this may be the cause of the above error.

Large Files

The problem can be caused by VMware VMFS3 limitation. With a default block size of 1 MB, this file system does not support files larger than 256 GB. VM disks larger than 256 GB will produce an error. To eliminate this problem, please migrate your source VM to a datastore which supports large files (e.g. VMFS5). If this problem occurs on the target datastore, use another datastore as a target.

Unexpected Server Outage During the Job Run

Source or Target host outage

During the job execution, the source or target host has been disconnected or powered off.

Please make sure the source host and/or the target host is powered on and accessible via network.

Source or Target Transporter outage

During the job execution, the source or target Transporter has been disconnected or powered off.

Please make sure the source host and/or the target host is powered on and accessible via network.

VM Quiescing Failure

Refer to the sections below to learn about possible causes of the following errors:

- “Failed to create quiesced snapshot of the VM”

If you have encountered this type of error, most likely VMware Tools or VSS is not properly configured on the source VM. To check if your environment is configured properly, please perform an easy check: Try to create a snapshot of the same source VM manually and select the "Quiesce guest file system" option. Please do this right after the NAKIVO Backup & Replication job fails so the load on the source VM is the same.

If the manual snapshot does not work, the problem is most likely caused by either VMware quiescing or Volume Shadow Copy Service. Please try the following:

- Update your virtual infrastructure (especially the hypervisor on the source host) to the latest version.
- Update VMware tools on the machines you protect to the latest version.
- Check if there are any Volume Shadow Copy Service errors in the Windows logs.
- After you resolve these issues, try to create another snapshot manually.

If you succeed, run the job in NAKIVO Backup & Replication and it should work as well.

Enabling CBT for a VM

Changed Block Tracking usage significantly speeds up incremental backup and replication.

Due to technical limitations, NAKIVO Backup & Replication cannot enable CBT automatically on Free ESXi hosts.

Free ESXi 5.1 and earlier: CBT can be enabled by editing the VM settings.

To enable CBT in a virtual machine which is running on Free ESXi 5.1 and earlier, follow the steps below:

1. Power off the virtual machine.
2. Right-click the virtual machine and click **Edit Settings**.
3. Click the **Options** tab.
4. Click **General** under the **Advanced** section and then click **Configuration Parameters**. The Configuration Parameters dialog opens.
5. Click **Add Row**.
6. Add the “*ctkEnabled*” parameter and then set its value to “*true*”.
7. Click **Add Row**, add “*scsi0:0.ctkEnabled*”, and set its value to “*true*”.

NOTE: The “*scsi0:0*” parameter in the “*scsi0:0.ctkEnabled*” value indicates the SCSI device assigned to the hard disk that is added to the virtual machine. Every hard disk added to the virtual machine is given a SCSI device that appears similar to *scsi0:0*, *scsi0:1*, or *scsi 1:1*.

8. Power on the virtual machine.

More information on Enabling CBT in VMware VMs can be found in VMware KB [1031873](#).

Free ESXi 5.5: CBT cannot be enabled without removing the VM from Inventory.

To enable CBT in a virtual machine which is running on Free ESXi 5.5, follow the steps below:

1. Power off the virtual machine.
2. Right-click the virtual machine and click **Remove from Inventory**.
3. Open Datastore Browser and browse the datastore where the virtual machine configuration file <VM name>.vmx is located.
4. Open the virtual machine folder on the datastore and download the <VM name>.vmx file.
5. Open the downloaded file using a text editor.
6. Manually add the *ctkEnabled* = “*true*” attribute.
7. Manually add the *scsi0:0.ctkEnabled* = “*true*” attribute.

NOTE: The “*scsi0:0*” parameter in the “*scsi0:0.ctkEnabled*” value indicates the SCSI device assigned to the hard disk that is added to the virtual machine. Every hard disk added to the virtual machine is given a SCSI device that appears similar to *scsi0:0*, *scsi0:1*, or *scsi 1:1*.

8. Save the modified file.
9. Upload the modified <VM name>.vmx file back to the datastore.
10. Right-click the uploaded file and click **Add to Inventory**. The wizard for adding to inventory opens.
11. Follow the wizard and complete adding the virtual machine to Inventory.
12. Power on the virtual machine.

Enabling SSH in ESXi Hosts

NAKIVO Backup & Replication supports Free ESXi hosts only with SSH access enabled. To enable SSH in ESXi host with vSphere Client, follow the steps below:

1. Select the host and click the **Configuration** tab.
2. Click **Security Profile > Properties**.
In ESXi 5.x, select **TSM-SSH, SSH** or **ESXi Shell** and click **Options**.
In ESXi 4.1, select **Local Tech Support** or **Remote Tech Support (SSH)** and click **Options**.
3. Choose the desired startup policy and click **Start**, then click **OK**.
We recommend choosing **Start and stop with host** as a startup policy.
4. Verify that the daemon selected in Step 3 shows as running in the Services Properties window.

More information on enabling SSH in ESXi host can be found in VMware KB [2004746](#).

Attributions

The following software components are used by NAKIVO Backup & Replication:

7zsfx

7-ZipLicense for use and distribution

7-Zip Copyright (C) 1999-2011 Igor Pavlov. Licenses for files are: 1) 7z.dll: GNU LGPL + unRAR restriction 2) All other files: GNU LGPL. The GNU LGPL + unRAR restriction means that you must follow both GNU LGPL rules and unRAR restriction rules.

Apache Commons Codec 1.6

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Apache HTTP Client 4.2.2

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Apache Mina 2.0.5

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Apache Tomcat 7.0.37

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

AspectJ 1.7.2

Eclipse Public License - v 1.0

<http://www.eclipse.org/legal/epl-v10.html>

AWS SDK for Java 1.5.5

Copyright 2010-2012 Amazon.com, Inc. or its affiliates. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. The Software shall be used for Good, not Evil.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Bcprov 1.46

Please note: our license is an adaptation of the MIT X11 License and should be read as such. Copyright (c) 2000 - 2011 The Legion Of The Bouncy Castle (<http://www.bouncycastle.org>)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Boost 1.53.0

Copyright Beman Dawes, David Abrahams, 1998-2005. Copyright Rene Rivera 2004-2007.

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the "Software") to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following: The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

C3PO 0.9

Licenses: [Eclipse Public License](#), [GNU Lesser General Public License](#).

Cglib 2.2.2

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Codehaus Jackson1.6.4

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Codehaus Jettison1.3

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons Codec 1.7

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons CLI 1.2

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons Collections 3.2

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons Email 1.3

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons Exec 1.1

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons FileUpload 1.3

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons IO 2.4

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Commons Lang 3.1

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Dom4j 1.6.6

Copyright (c) 2001-2005 MetaStuff, Ltd. All Rights Reserved.

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain copyright statements and notices. Redistributions must also contain a copy of this document. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. The name "DOM4J" must not be used to endorse or promote products derived from this Software without prior written permission of MetaStuff, Ltd. For written permission, please contact dom4j-info@metastuff.com. 4. Products derived from this Software may not be called "DOM4J" nor may "DOM4J" appear in their names without prior written permission of MetaStuff, Ltd.

DOM4J is a registered trademark of MetaStuff, Ltd. 5. Due credit should be given to the DOM4J Project - <http://dom4j.sourceforge.net>

THIS SOFTWARE IS PROVIDED BY METASTUFF, LTD. AND CONTRIBUTORS ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL METASTUFF, LTD. OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Ehcache 2.6

Copyright 2003-2010 Terracotta, Inc.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Elevate.exe

Copyright © 2007 Free Software Foundation, Inc.

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version. This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

ExtJS 4.0.7

Ext JS - JavaScript Library

Copyright (c) 2006-2011, Sencha Inc.

All rights reserved. <http://www.sencha.com/license>

Fast MD5 2.7.1

Copyright (C) <2012> Free Software Foundation

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version. This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

Flying Saucer R8

Copyright © 2007 Free Software Foundation, Inc.

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 3 of the License, or (at your option) any later version. This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

Fugue Icons 3.3

Copyright (c) 2012 Massachusetts Institute of Technology

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Glib 2.16.4

GNU LIBRARY GENERAL PUBLIC LICENSE Version 2, June 1991 Copyright (C) 1991 Free Software Foundation, Inc.
<http://www.gnu.org/licenses/old-licenses/lgpl-2.0.html>

Google Gson 1.7.1

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Google Guava r09

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS,

WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Google protobuf 2.4.1

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

H2 database

Eclipse Public License - v 1.0

<http://www.eclipse.org/legal/epl-v10.html>

Hibernate 4.2

Copyright (C) 2012 Free Software Foundation, Inc.

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

ICU4J 51

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1995-2014 International Business Machines Corporation and others

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice

appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

All trademarks and registered trademarks mentioned herein are the property of their respective owners..

Jadira Usertype 3.1

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Java Image Scaling 0.8.5

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Javamail 1.4

The JavaMail reference implementation is licensed under the Common Development and Distribution License (CDDL) v1.1 and GNU General Public License (GPL) v2 with Classpath Exception.

COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL) Version 1.1

1. Definitions.

- 1.1. "Contributor" means each individual or entity that creates or contributes to the creation of Modifications.
- 1.2. "Contributor Version" means the combination of the Original Software, prior Modifications used by a Contributor (if any), and the Modifications made by that particular Contributor.
- 1.3. "Covered Software" means (a) the Original Software, or (b) Modifications, or (c) the combination of files containing Original Software with files containing Modifications, in each case including portions thereof.
- 1.4. "Executable" means the Covered Software in any form other than Source Code.
- 1.5. "Initial Developer" means the individual or entity that first makes Original Software available under this License.
- 1.6. "Larger Work" means a work which combines Covered Software or portions thereof with code not governed by the terms of this License.
- 1.7. "License" means this document.
- 1.8. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.
- 1.9. "Modifications" means the Source Code and Executable form of any of the following:
 - A. Any file that results from an addition to, deletion from or modification of the contents of a file containing Original Software or previous Modifications;
 - B. Any new file that contains any part of the Original Software or previous Modification; or
 - C. Any new file that is contributed or otherwise made available under the terms of this License.
- 1.10. "Original Software" means the Source Code and Executable form of computer software code that is originally released under this License.
- 1.11. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.
- 1.12. "Source Code" means (a) the common form of computer software code in which modifications are made and (b) associated documentation included in or with such code.
- 1.13. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. License Grants.

2.1. The Initial Developer Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, the Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license:

- (a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer, to use, reproduce, modify, display, perform, sublicense and distribute the Original Software (or portions thereof), with or without Modifications, and/or as part of a Larger Work; and
- (b) under Patent Claims infringed by the making, using or selling of Original Software, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Software (or portions thereof).
- (c) The licenses granted in Sections 2.1(a) and (b) are effective on the date Initial Developer first distributes or otherwise makes the Original Software available to a third party under the terms of this License.
- (d) Notwithstanding Section 2.1(b) above, no patent license is granted: (1) for code that You delete from the Original Software, or (2) for infringements caused by: (i) the modification of the Original Software, or (ii) the combination of the Original Software with other software or devices.

2.2. Contributor Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

- (a) under intellectual property rights (other than patent or trademark) Licensable by Contributor to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof), either on an unmodified basis, with other Modifications, as Covered Software and/or as part of a Larger Work; and
- (b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: (1) Modifications made by that Contributor (or portions thereof); and (2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).
- (c) The licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first distributes or otherwise makes the Modifications available to a third party.
- (d) Notwithstanding Section 2.2(b) above, no patent license is granted: (1) for any code that Contributor has deleted from the Contributor Version; (2) for infringements caused by: (i) third party modifications of Contributor Version, or (ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or (3) under Patent Claims infringed by Covered Software in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Availability of Source Code.

Any Covered Software that You distribute or otherwise make available in Executable form must also be made available in Source Code form and that Source Code form must be distributed only under the terms of this License. You must include a copy of this License with every copy of the Source Code form of the Covered Software You distribute or otherwise make available. You must inform recipients of any such Covered Software in Executable form as to how they can obtain such Covered Software in Source Code form in a reasonable manner on or through a medium customarily used for software exchange.

3.2. Modifications.

The Modifications that You create or to which You contribute are governed by the terms of this License. You represent that You believe Your Modifications are Your original creation(s) and/or You have sufficient rights to grant the rights conveyed by this License.

3.3. Required Notices.

You must include a notice in each of Your Modifications that identifies You as the Contributor of the Modification. You may not remove or alter any copyright, patent or trademark notices contained within the Covered Software, or any notices of licensing or any descriptive text giving attribution to any Contributor or the Initial Developer.

3.4. Application of Additional Terms.

You may not offer or impose any terms on any Covered Software in Source Code form that alters or restricts the applicable version of this License or the recipients' rights hereunder. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, you may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.5. Distribution of Executable Versions.

You may distribute the Executable form of the Covered Software under the terms of this License or under the terms of a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable form does not attempt to limit or alter the recipient's rights in the Source Code form from the rights set forth in this License. If You distribute the Covered Software in Executable form under a different license, You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.6. Larger Works.

You may create a Larger Work by combining Covered Software with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Software.

4. Versions of the License.

4.1. New Versions.

Oracle is the initial license steward and may publish revised and/or new versions of this License from time to time. Each version will be given a distinguishing version number. Except as provided in Section 4.3, no one other than the license steward has the right to modify this License.

4.2. Effect of New Versions.

You may always continue to use, distribute or otherwise make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. If the Initial Developer includes a notice in the Original Software prohibiting it from being distributed or otherwise made available under any subsequent version of the License, You must distribute and make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. Otherwise, You may also choose to use, distribute or otherwise make the Covered Software available under the terms of any subsequent version of the License published by the license steward.

4.3. Modified Versions.

When You are an Initial Developer and You want to create a new license for Your Original Software, You may create and use a modified version of this License if You: (a) rename the license and remove any references to the name of the license steward (except to note that the license differs from this License); and (b) otherwise make it clear that the license contains terms which differ from this License.

5. DISCLAIMER OF WARRANTY.

COVERED SOFTWARE IS PROVIDED UNDER THIS LICENSE ON AN “AS IS” BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED SOFTWARE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED SOFTWARE IS WITH YOU. SHOULD ANY COVERED SOFTWARE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

6. TERMINATION.

6.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

6.2. If You assert a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You assert such claim is referred to as “Participant”) alleging that the Participant Software (meaning the Contributor Version where the Participant is a Contributor or the Original Software where the Participant is the Initial Developer) directly or indirectly infringes any patent, then any and all rights granted directly or indirectly to You by such Participant, the Initial Developer (if the Initial Developer is not the Participant) and all Contributors under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively and automatically at the expiration of such 60 day notice period, unless if within such 60 day period You withdraw Your claim with respect to the Participant Software against such Participant either unilaterally or pursuant to a written agreement with Participant.

6.3. If You assert a patent infringement claim against Participant alleging that the Participant Software directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

6.4. In the event of termination under Sections 6.1 or 6.2 above, all end user licenses that have been validly granted by You or any distributor hereunder prior to termination (excluding licenses granted to You by any distributor) shall survive termination.

7. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED SOFTWARE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

8. U.S. GOVERNMENT END USERS.

The Covered Software is a “commercial item,” as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of “commercial computer software” (as that term is defined at 48 C.F.R. § 252.227-7014(a)(1)) and “commercial computer software documentation” as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Software with only those rights set forth herein. This U.S. Government Rights clause is in lieu of, and supersedes, any other FAR, DFAR, or other clause or provision that addresses Government rights in computer software under this License.

9. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by the law of the jurisdiction specified in a notice contained within the Original Software (except to the extent applicable law, if any, provides otherwise), excluding such jurisdiction's conflict-of-law provisions. Any litigation relating to this License shall be subject to the jurisdiction of the courts located in the jurisdiction and venue specified in a notice contained within the Original Software, with the losing party responsible for costs, including, without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License. You agree that You alone are responsible for compliance with the United States export administration regulations (and the export control laws and regulation of any other countries) when You use, distribute or otherwise make available any Covered Software.

10. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

NOTICE PURSUANT TO SECTION 9 OF THE COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)

The code released under the CDDL shall be governed by the laws of the State of California (excluding conflict-of-law provisions). Any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California and the state courts of the State of California, with venue lying in Santa Clara County, California.

The GNU General Public License (GPL) Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software

Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

"CLASSPATH" EXCEPTION TO THE GPL VERSION 2

Certain source files distributed by Oracle are subject to the following clarification and special exception to the GPL Version 2, but only where Oracle has expressly included in the particular source file's header the words "Oracle designates this particular file as subject to the "Classpath" exception as provided by Oracle in the License file that accompanied this code."

Linking this library statically or dynamically with other modules is making a combined work based on this library. Thus, the terms and conditions of the GNU General Public License Version 2 cover the whole combination.

As a special exception, the copyright holders of this library give you permission to link this library with independent modules to produce an executable, regardless of the license terms of these independent modules, and to copy and distribute the resulting executable under terms of your choice, provided that you also meet, for each linked independent module, the terms and conditions of the license of that module. An independent module is a module which is not derived from or based on this library. If you modify this library, you may extend this exception to your version of the library, but you are not obligated to do so. If you do not wish to do so, delete this exception statement from your version.

Jaxen 1.1.1

Copyright 2003-2006 The Werken Company.

All Rights Reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided

with the distribution. 3. Neither the name of the Jaxen Project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

JNA JAR 3.4.0

Copyright (C) 2004 ASF

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

Joda Time 2.2

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

JRE 1.7.0_51

NAKIVO Backup & Replication does not use any Commercial Features of Java SE Product Editions.

JSoup 1.7.2

Copyright © 2009 - 2013 Jonathan Hedley

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM,

DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Libsedb

GNU LESSER GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

0. Additional Definitions.

As used herein, “this License” refers to version 3 of the GNU Lesser General Public License, and the “GNU GPL” refers to version 3 of the GNU General Public License.

“The Library” refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An “Application” is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A “Combined Work” is a work produced by combining or linking an Application with the Library. The particular version of the Library with which the Combined Work was made is also called the “Linked Version”.

The “Minimal Corresponding Source” for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The “Corresponding Application Code” for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

1. Exception to Section 3 of the GNU GPL.

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

2. Conveying Modified Versions.

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

- a) under this License, provided that you make a good faith effort to ensure that, in the event an Application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or
- b) under the GNU GPL, with none of the additional permissions of this License applicable to that copy.

3. Object Code Incorporating Material from Library Header Files.

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

- a) Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the object code with a copy of the GNU GPL and this license document.

4. Combined Works.

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

- a) Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the Combined Work with a copy of the GNU GPL and this license document.
- c) For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.
- d) Do one of the following:
 - 0) Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.
 - 1) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (a) uses at run time a copy of the Library already present on the user's computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.
- e) Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 4d0, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option 4d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

5. Combined Libraries.

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.
- b) Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

6. Revised Versions of the GNU Lesser General Public License.

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy's public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.

Liquibase 2.0

Copyright (C) 2006-2014 Nathan Voxland

This library is free software; you can redistribute it and/or modify it under the terms of the Apache License, Version 2.0.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

Log4j 2.0

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Lookup3.c

lookup3.c, by Bob Jenkins, May 2006, Public Domain.

OneJar 0.97

One-JAR™ (<http://www.simontuffs.com/one-jar>). Copyright (c) 2004-2010,

P. Simon Tuffs (simon@simontuffs.com). All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of P. Simon Tuffs, nor the names of any contributors, nor the name One-JAR may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Including this file inside the built One-JAR file conforms with these terms.

Orion SSH 2

Copyright (c) 2007-2008 Trilead AG (<http://www.trilead.com>)

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- a.) Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- b.) Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- c.) Neither the name of Trilead nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Trilead SSH-2 for Java includes code that was written by Dr. Christian Plattner during his PhD at ETH Zurich. The license states the following:

Copyright (c) 2005 - 2006 Swiss Federal Institute of Technology (ETH Zurich),

Department of Computer Science (<http://www.inf.ethz.ch>),

Christian Plattner. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- a.) Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- b.) Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

c.) Neither the name of ETH Zurich nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The Java implementations of the AES, Blowfish and 3DES ciphers have been taken (and slightly modified) from the cryptography package released by "The Legion Of The Bouncy Castle". Their license states the following:

Copyright (c) 2000 - 2004 The Legion Of The Bouncy Castle

(<http://www.bouncycastle.org>)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

OpenSSL 0.9.8x

OpenSSL License

Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

<http://www.openssl.org/source/license.html>

OpenSSL 0.9.8y

OpenSSL License

Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

<http://www.openssl.org/source/license.html>

Pkware ZIP Appnote 6.3.3

Copyright (c) 1989 - 2012 PKWARE Inc., All Rights Reserved.

<http://www.pkware.com/documents/casestudies/APPNOTE.TXT>

Although PKWARE will attempt to supply current and accurate information relating to its file formats, algorithms, and the subject programs, the possibility of error or omission cannot be eliminated. PKWARE therefore expressly disclaims any warranty that the information contained in the associated materials relating to the subject programs and/or the format of the files created or accessed by the subject programs and/or the algorithms used by the subject programs, or any other matter, is current, correct or accurate as delivered. Any risk of damage due to any possible inaccurate information is assumed by the user of the information. Furthermore, the information relating to the subject programs and/or the file formats created or accessed by the subject programs and/or the algorithms used by the subject programs is subject to change without notice.

Postgresql connector 9.0-801.jdbc4

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

SLF4J 1.7.5

Copyright (c) 2012 Massachusetts Institute of Technology

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Snappy 1.1.1

Copyright 2011, Google Inc.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided

with the distribution. * Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Spring Framework 3.2

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

SWT

Eclipse Public License - v 1.0

<http://www.eclipse.org/legal/epl-v10.html>

Thymeleaf 2.0.14

Copyright 2004 ASF

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

VI Java 5.5

Copyright (c) 1998, Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR

CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

VMware VDDK 5.1.1

Copyright 2007-2010 VMware, Inc. All rights reserved.

This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware, the VMware "boxes" logo and design, Virtual SMP, and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware VDDK 5.5

© 1998-2013 VMware, Inc. All rights reserved. This product is protected by copyright and intellectual property laws in the United States and other countries as well as by international treaties. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Libesedb 1.0.0 (20120102)

GNU LESSER GENERAL PUBLIC LICENSE Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>

<http://www.gnu.org/licenses/lgpl-3.0.html>

Libuna 1.0.0 (20140105)

GNU LESSER GENERAL PUBLIC LICENSE Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>

<http://www.gnu.org/licenses/lgpl-3.0.html>

WinRun4J 0.4.4

Common Public License Version 1.0 License agreement <http://www.opensource.org/licenses/cpl1.0>

XStream 1.2.2

Copyright ©2012 The FreeBSD Project. All Rights Reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY [LICENSOR] "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Zlib 1.2.5/1.2.3

Copyright © 1995-2012 Jean-loup Gailly and Mark Adler.

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software. Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions: 1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required. 2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software. 3. This notice may not be removed or altered from any source distribution.